PingFederate Server



Contents

Release Notes	Ç
PingFederate 10.1.11 - February 2023	
PingFederate 10.1.10 - January 2022	
PingFederate 10.1.9 - October 2021	
PingFederate 10.1.8 - August 2021	
PingFederate 10.1.7 - June 2021	
PingFederate 10.1.6 - May 2021	
PingFederate 10.1.5 - March 2021	
PingFederate 10.1.4 - January 2021	
PingFederate 10.1.3 - December 2020	
PingFederate 10.1.2 - October 2020	
PingFederate 10.1.1 - August 2020	
PingFederate 10.1 - June 2020	
Navigation tabs and menus	
Known issues and limitations	
Deprecated features	
Previous releases	
11001040 10104000	
Introduction to PingFederate	27
About identity federation and SSO	
Service providers and identity providers	
Federation hub	
Supported standards	
Federation roles	
Terminology	
Browser-based SSO	
Web services standards	
OAuth 2.0	
System for Cross-domain Identity Management (SCIM)	
Transport and message security	
SSO integration overview	
SSO integration concepts	
Identity provider integration	
Service provider integration	
Integrations and deployment scenarios	
Security token service	
OAuth authorization server	
User account management	
Enterprise deployment architecture	
Additional features	
Key concepts	
About WS-Trust STS	
About OAuth	
SSO integration kits and adapters	
Security infrastructure	
Hierarchical plugin configurations	
Identity mapping	
User attributes	

User provisioning	97
Customer identity and access management	
Federation hub use cases	
Federation planning checklist	
	400
Installing PingFederate	
System requirements	
Database driver information	
Port requirements	
Installing Java	
Installation options	
Installing PingFederate on Windows	
Installing PingFederate on Linux systems	
Installing the PingFederate service on Linux manually Installing PingFederate service on Windows manually	
motaling i mg. odorato corrido en vindono mandany	
Uninstalling PingFederate	126
Uninstalling PingFederate from a Windows server	
Uninstalling PingFederate from a Linux server	
Upgrading PingFederate	128
Upgrade considerations	
Upgrade considerations introduced in PingFederate 9.x	
Upgrade considerations introduced in PingFederate 8.x	
Upgrade considerations introduced in PingFederate 7.x	
Upgrade considerations introduced in PingFederate 6.x	
Updating to the latest maintenance release	
Upgrading PingFederate on Windows using the installer	137
Upgrading PingFederate on Windows using the Upgrade Utility	
Upgrading PingFederate on Linux systems	
Custom mode	
Reviewing post-upgrade tasks	
Reviewing administrative users	
Copying customized files or settings	
Reviewing database changes	
Reviewing log configuration	
Migrating other components	
Resetting files and variable for HSMVerifying the new installation	
Vollying the new installation	
Getting Started with PingFederate	154
Starting and stopping PingFederate servers	
Starting and stopping PingFederate on Windows	
Starting and stopping PingFederate on Linux	
Opening the PingFederate administrative console	
Setting up PingFederate	
Importing your license	
Entering basic information for PingFederate	
Configuring identity provider settings	
Creating an administrator account	
Reviewing the PingFederate configuration	
Setting up PingFederate Bridge	161

	Connecting PingFederate to PingOne for Enterprise	161
	Connecting PingFederate Bridge to a directory server	161
	Configuring PingOne SSO and PingID VPN settings	162
	Entering basic information for PingFederate Bridge	
	Reviewing the PingFederate Bridge configuration	
	Completion of the PingFederate Bridge setup	166
	PingFederate administrative console	
	Tasks and steps	
	Console buttons	
	Third-party cryptographic solutions	
	Supported hardware security modules	
	Supported software security package	178
Se	erver Clustering Guide	182
	Overview of clustering	
	Cluster protocol architecture	
	Runtime state-management architectures	
	Adaptive clustering	
	Directed clustering	
	Runtime state-management services	
	Inter-Request State-Management (IRSM) Service	
	IdP Session Registry Service	
	SP Session Registry Service	
	LRU memory management schemes	
	Assertion Replay Prevention Service	
	Artifact-Message Persistence and Retrieval Service	
	Back-Channel Session Revocation Service	
	Account Locking Service	
	Other services	
	Deploying cluster servers	
	Enabling dynamic discovery for clustering	
	Deploying provisioning failover	
	Configuration synchronization	
	Console configuration push	213
	Configuration-archive deployment	213
۸،	dministrator's Reference Guide	21.4
Α(
	Attribute mapping expressions	
	Enabling and disabling expressions	
	Construct OGNL expressions	
	Using the OGNL edit window	
	Authentication policies	
	Policies	
	Policy contractsAdapter Mappings	
	Sessions	
	Bundled adapters	
	Composite Adapter	
	HTML Form Adapter	
	HTTP Basic Adapter	
	Identifier First Adapter	
	Kerberos Adapter	
	OpenTakan Adapter	210

Customer IAM configuration	324
Setting up PingDirectory for customer identities	
Managing local identity profiles	
Configuring the HTML Form Adapter for customer identities	
Setting up self-service registration	
Enabling third-party identity providers without registration	
Customizing assertions and authentication requests	
Message types and available variables	
Sample customizations	
Fulfillment by datastore queries	
Attribute mapping with multiple data sources	362
Datastore query configuration	362
IdP-to-SP bridging	371
Adapter-to-adapter mappings	
Token translator mappings	377
Identity provider SSO configuration	
IdP application integration settings	
IdP protocol endpoints	
SP connection management	
SP affiliations	
OAuth configuration	
Configuring OAuth use cases	
Configuring authorization server settings	
Scopes and scope management	
Configuring client settings	
Managing Client Registration Policy instances	
Managing OAuth clients	
Grant contract mapping	
Token mapping	
Client Initiated Backchannel Authentication (CIBA)	
OAuth attribute mapping using a datastore	592
OAuth client session management	593
OAuth token exchange	594
Security management	599
Certificate and key management	600
System integration	631
Account lockout protection	637
Password spraying prevention	
Implementing a MasterKeyEncryptor using AWS KMS	
Self-service user account management	
Configuring self-service password management	
Configuring self-service account recovery	
Configuring self-service user name recovery	
Service provider SSO configuration	
SP application integration settings	
Federation settings	
Managing IdP connections	
OpenID Connect Relying Party support	
Configuring IdP discovery using a persistent cookie	
System administration	
Configuring PingFederate properties	
Enabling OIDC-based authentication	
Configuring size limits	
PingFederate log files	
Alternative console authentication	780
Configuring automatic connection validation	790

Automating configuration migration	790
Outbound provisioning CLI	
Customizable user-facing pages	798
Customizable email notifications	810
Customizable text message	817
Localizing messages for end users	817
Configuring a password policy	
Managing cipher suites	
Manage externally stored authentication sessions	
OAuth persistent grants cleanup	
Specifying the domain of the PF cookie	
Specifying the domain of the PF.PERSISTENT cookie	
Extending the lifetime of the PingFederate cookie	
Configuring forward proxy server settings	
Adding custom HTTP response headers	
Configuring validation for the AudienceRestriction element	
Customizing the OpenID Provider configuration endpoint response	
Customizing the heartbeat message	
Customizing the favicon for application and protocol endpoints	
Configuring the behavior of searching multiple datastores with one mapping	
System settings	
Server	
Metadata	858
Monitoring and notifications	863
Datastores	867
Password Credential Validators	905
Active Directory and Kerberos	915
External systems	919
Troubleshooting	933
Enabling debug messages and console logging	933
Resolving startup issues	936
Troubleshooting data store issues	936
Resolving URL-related errors	936
Resolving service-related errors	937
Troubleshooting authentication policy issues	938
Troubleshooting registration and profile management issues	940
Troubleshooting runtime errors	
Troubleshooting OAuth transactions	944
Other runtime issues	949
Collecting support data	
WS-Trust STS configuration	
Server settings	
Identity provider STS configuration	
Service provider STS configuration	975
Double was a Trusting Codd to	000
Performance Tuning Guide	
Logging	
Operating system tuning	
Linux tuning	
Windows tuning	
Concurrency	
Tuning the acceptor queue size	
Tuning the server thread pool	
Configuring connection pools to datastores	
Memory	994

	994
•	995
	996
• • • • • • • • • • • • • • • • • • • •	997
	1005
•	1005
References	1005
PingFederate Monitoring Guide	1005
	1006
·	1008
	1008
	1016
	1017
Splunk dashboards and audit logs	1018
•	1021
	1023
	1024
	1024
	1025
	1029
	1031
	1031
	tors
· · · · · · · · · · · · · · · · · · ·	1034
	1041
	1042
Log messages	
Developer's Reference Guide	1043
•	1043
·	1043
	tion endpoint
	1057
•	
·	
	1071
	1073
·	1080
	nt1081
·	
	1086
	s1092
The second secon	

OAuth Client Management Service	1093
OAuth Access Grant Management Service	1114
OAuth Persistent Grant Management API	1117
Session Management API	1120
Session Revocation API endpoint	
PingFederate administrative API	1125
Application endpoints	
IdP endpoints	
SP endpoints	
System-services endpoints	
Authentication API	
Exploring the authentication API	
Mobile application authentication through REST APIs	
Development of authentication API-capable adapters and selectors	
Authentication API states, actions, and models	
Specification of the plugin API	
State model contents	
Non-interactive plugins	
Runtime behavior implementation	
Session state management	
Error messages and localization	1169
Legal Information	1170
Index	1171

Release Notes

These release notes summarize the changes in current and previous product updates.

PingFederate enables outbound and inbound solutions for single sign-on (SSO), federated identity management, mobile identity security, API security, social identity integration, and customer identity and access management. PingFederate extends employee, customer, and partner identities across domains without passwords, using only standard identity protocols: SAML, WS-Federation, WS-Trust, OAuth, and SCIM.

PingFederate 10.1.11 - February 2023

PingFederate 10.1.11 is a cumulative maintenance release for PingFederate 10.1. For a summary of the features introduced in the 10.1 release, see *PingFederate 10.1 - June 2020* on page 15.

Resolved issues

Ticket ID	Description
PF-32805	We've resolved a potential security vulnerability that is described in security advisory SECADV033.

PingFederate 10.1.10 - January 2022

PingFederate 10.1.10 is a cumulative maintenance release for PingFederate 10.1. For a summary of the features introduced in the 10.1 release, see *PingFederate 10.1 - June 2020* on page 15.

Ticket ID	Description
PF-30052	Added debug-level server logs to show lifecycle information for internally managed access tokens that access token managers generate.
PF-30450	Resolved a potential security vulnerability that is described in security bulletin SECBL021.
PF-30536	Resolved a potential security vulnerability by updating Apache Log4j2 to version 2.17.1.

PingFederate 10.1.9 is a cumulative maintenance release for PingFederate 10.1. For a summary of the features introduced in the 10.1 release, see *PingFederate 10.1 - June 2020* on page 15.

Resolved issues

Ticket ID	Description
PF-29512	Resolved a harmless log entry that appeared as an org.apache.xml.security.Init error when you updated PingFederate to versions 9.2.3-13 or 10.0.11.
PF-29621	In the in-place upgrade <code>.zip</code> package for maintenance releases, configurable files in the <code>pingfederate/bin</code> folder are now placed under the <code>pf-maintenance/merge_required</code> folder to indicate that you must merge these updated files with your existing copies.
PF-29924	Resolved a potential security vulnerability involving authentication policies.

PingFederate 10.1.8 - August 2021

PingFederate 10.1.8 is a cumulative maintenance release for PingFederate 10.1. For a summary of the features introduced in the 10.1 release, see *PingFederate 10.1 - June 2020* on page 15.

Resolved issues

Ticket ID	Description
PF-29310	Resolved a potential security vulnerability.
PF-29366	Resolved a potential security vulnerability.

PingFederate 10.1.7 - June 2021

PingFederate 10.1.7 is a cumulative maintenance release for PingFederate 10.1. For a summary of the features introduced in the 10.1 release, see *PingFederate 10.1 - June 2020* on page 15.

Ticket ID	Description
PF-28831	Resolved a potential security vulnerability involving the authentication API.
PF-28846	Enhanced security by no longer allowing the PingFederate web service to serve the files contained in <pf_install>/pingfederate/server/default/conf/template.</pf_install>

PingFederate 10.1.6 is a cumulative maintenance release for PingFederate 10.1. For a summary of the features introduced in the 10.1 release, see *PingFederate 10.1 - June 2020* on page 15.

Resolved issues

Ticket ID	Description		
PF-28564	Updated jackson-databind to version 2.9.10.8.		
PF-28683	Updated the Apache Velocity engine with security patches.		
PF-28729	Resolved an issue that prevented PingFederate from using the OAuth device authorization grant type when authenticating users with an IdP connection.		
PF-28731	Provided a workaround for an issue that prevents administrators running PingFederate 10.1.0 from performing a rolling upgrade to a later 10.1 maintenance release. To enable the workaround, set rebalance.lock.required=false in the cluster-adaptive.conf file on the new nodes before startup. After completing the upgrade, set this property to true or remove it.		

PingFederate 10.1.5 - March 2021

PingFederate 10.1.5 is a cumulative maintenance release for PingFederate 10.1. For a summary of the features introduced in the 10.1 release, see *PingFederate 10.1 - June 2020* on page 15.

Ticket ID	Description		
PF-28423	Resolved an issue that caused the audit log to intermittently record incorrect connection IDs.		
PF-28433	To reduce the number of search requests to directory servers during user authentication, PingFederate no longer performs multiple calls to the LDAP Root DSE checking for IntermediateClientRequestControl support. This LDAP control can send servers information about a client and any downstream clients.		
PF-28435	The OAuth client connection ID is no longer missing from some events in the audit log.		
PF-28464	PingFederate with Bouncy Castle FIPS integration now lets you place the Bouncy Castle provider at the top of the provider list in the java.security file. This ensures that PingFederate uses the Bouncy Castle provider for cryptographic signing operations. Before you change the provider ordering in an existing deployment, see "Integrating with Bouncy Castle FIPS provider" in the PingFederate documentation.		

PingFederate 10.1.4 - January 2021

PingFederate 10.1.4 is a cumulative maintenance release for PingFederate 10.1. For a summary of the features introduced in the 10.1 release, see *PingFederate 10.1 - June 2020* on page 15.

Resolved issues

Ticket ID	Description
PF-27095	Resolved an issue that caused the OAuth bearer access token token processor to ignore issuer and audience values in access tokens.
PF-28047	Resolved a potential security vulnerability described in security bulletin SECBL018 on the <i>Ping Identity Support</i> website.

PingFederate 10.1.3 - December 2020

PingFederate 10.1.3 is a cumulative maintenance release for PingFederate 10.1. For a summary of the features introduced in the 10.1 release, see *PingFederate 10.1 - June 2020* on page 15.

Ticket ID	Description		
PF-27620	Fixed an issue that prevented the selection of the expressions source type for authentication policy contract mappings that use datastores in SP connection administrative API requests.		
PF-27650	SSL server certificates and SSL certificate signing requests generated by PingFederate now comply with macOS Catalina's new restrictions, so browsers using macOS's trust store no longer fail.		
PF-27678	Fixed support for unauthenticated SMTP, allowing PingFederate notification publishers to connect to SMTP servers when the usernames or passwords are blank.		
PF-27723	Fixed an issue that could increase heap usage for deployments with frequent federation hub SSO workflows because the transaction state was not purged from memory on completion of the workflow.		
PF-27740	When users changing or resetting their password enter a password that does not meet the policy requirements, now the form shows the password requirements instead of an error message. This issue occurred only if PingFederate was connected to PingDirectory 8.x and PingFederate enabled "PingDirectory Detailed Password Policy Requirement Messaging" on an LDAP username password credential validator.		
PF-27749	Improved performance for persistent session LDAP storage.		

PingFederate 10.1.2 is a cumulative maintenance release for PingFederate 10.1. For a summary of the features introduced in the 10.1 release, see *PingFederate 10.1 - June 2020* on page 15.

Ticket ID	Description		
PF-27043	Resolved an issue that could cause a "file does not exist" error to impact user requests when cluster replication is performed while the cluster is under load.		
PF-27067	Resolved an issue causing the HTML Form Adapter not to display the custom template that is supposed to be displayed when a user clicks Cancel during the password reset flow.		
PF-27076	Updated jackson-databind to version 2.9.10.6.		
PF-27077	Resolved an issue preventing the configurations for OAuth clients created by Dynamic Client Registration to be updated after PingFederate was upgraded. An error was displayed stating, "An updated client secret is required because one of the chosen algorithms relies on a shared symmetric secret."		
PF-27097	When PingFederate sends a LDAP transaction request to a PingDirectory server, if the request specifies an intermediate client request control, then PingFederate no longer specifies that the control is critical. This change prevents runtime errors when the directory server does not support the control but is behind a proxy server that does support it.		
PF-27105	Resolved an issue where some administrative console pages did not render when using Luna in FIPS 140-2 approved operation mode.		
PF-27297	When an OAuth client uses the redirectless mode to initiate authorization directly through the authentication API, PingFederate no longer returns an error if the client has multiple redirect URIs or has a redirect URI with a wildcard character.		
PF-27337	Enhanced certification revocation list (CRL) processing to prevent concurrent calls to the same CRL endpoint.		
PF-27356	PingFederate no longer relies on an externally-hosted CSS library for font styles.		
PF-27385	When PingFederate cannot find a grant, now it logs that it did not find the grant instead of logging a null pointer exception and stack trace.		
PF-27422	Improved logging for signed request objects that PingFederate ignores and added a config-store flag to always process request objects.		
PF-27463	PingFederate now uses defaultForLegacyConfig from the field descriptor to fill in missing field values in POST and PUT requests to the administrative API for plugins. This reduces the chance of validation errors when a new configuration field has been added to the plugin and a POST or PUT request omits the field.		
PF-27504	Resolved an issue that caused the administrative console to fail when PingFederate is configured to use an external datastore for OAuth client storage but the datastore is not available.		
PF-27541	PingFederate now respects the PKCE authorization request parameters from the request object.		

PingFederate 10.1.1 is a cumulative maintenance release for PingFederate 10.1. For a full summary of the features introduced in the 10.1 release, see *PingFederate 10.1 - June 2020* on page 15.

Description		
Added a configurable field called Trim username spaces for search to the LDAP password credential validator. When the field is enabled, leading and trailing spaces are removed from the username in the search filter and in the returned username attribute. When creating new LDAP password credential validators, this field is enabled by default.		
When PingFederate is configured with non-native administrative console authentication, the engine nodes in a cluster no longer try to authenticate to the backend service using the mode configured for the console node.		
When a user is a member of multiple provisioning groups managed by different channels, removing the user from the group that provisioned them will no longer immediately deprovision the user. Instead, when the provisioning cycle runs, PingFederate re-assigns the user to a different channel. PingFederate might update the user attributes if the new channel has a different attribute mapping.		
During a device authorization flow, PingFederate now returns an error when an OAuth client uses a client_id to get an access token that does not match the client_id it used to get the device_code.		
The administrative API and OAuth client management service now return only current client lists.		
The authentication API redirectless flow now supports self-service password reset with authentication policies.		
During a user authentication or password change, if CAPTCHA is enabled and the CAPTCHA response is empty, PingFederate now shows a CAPTCHA error instead of a null pointer exception.		
he OAuth access grant management service can now GET and DELETE grants ssigned to disabled OAuth clients.		
Resolved a PingFederate 10.1 issue that affected the performance of authentication and ession creation in clustered environments.		
For OAuth and OIDC administrative authentication, PingFederate now validates the endpoints in the OAuth and OIDC properties files to ensure that they are HTTPS. However, the OIDC end.session.endpoint still does not require HTTPS.		
Updated jackson-databind to version 2.9.10.5.		
The pfbrowserid cookie is now set with the secure and httponly flags.		
PingFederate 10.1 templates now load correctly when you have a CIBA authenticator configured with the PingID SDK adapter 1.6.		
After you import a certificate signing response into a certificate, PingFederate now persists the updated certificate.		
mproved PingFederate's performance when processing scopes and scope groups during and user requests.		

Ticket ID	Description	
PF-26651	Rescinded the requirement, introduced in PingFederate 10.1, to obfuscate the values of client secrets used for OIDC and OAuth administrative authentication. Instead, now when PingFederate reads a plain text value for an OIDC, OAuth, LDAP, or RADIUS shared secret property, it logs a WARN level message that the value should be obfuscated.	
PF-26659	On the administrative console, when you change settings and then click Save , now the console shows a message stating whether PingFederate saved the new settings. Also, resolved some minor UI issues.	
PF-26678	Updated the Apache Commons IO library in the PingFederate Upgrade Utility to version 2.5.	
PF-26796	The IdP-discovery common domain cookie is now set with the secure and httponly flags.	
PF-26797	Resolved an issue that caused the Show Me Around tutorial to appear each time the use logs in to PingFederate, even if the user had finished the tutorial, or had clicked Dismiss or the X icon on the tutorial's popup.	
PF-26828	On the Authorization server settings window, the CORS Allowed Origins setting now supports non-HTTP/HTTPS formatted values (registered schemes).	
PF-26844	Resolved an issue introduced in PingFederate 10.1 that prevented it from correctly storing the user ID of the IdP connection's authentication source in the pf-connected-identity attribute.	
PF-26853	Resolved a potential security vulnerability described in security bulletin SECBL017 on the <i>Ping Identity Support</i> website.	
PF-26868	Improved logging surrounding access grant manager JDBC interactions.	

PingFederate 10.1 - June 2020

Summary of the latest enhancements, resolved issues, and other improvements for PingFederate.

Enhancements

Administrative console and API

A new delightful experience

The use cases and standards that PingFederate supports have expanded over the years. Organizations implementing workforce and customer solutions are increasingly playing multiple authentication and federation roles. As a result, we are modernizing the PingFederate administrative console. The new dashboard provides **Shortcuts** to common configurations and **Helpful Links** that guide you and your PingFederate administrators to complete tasks faster. You can name your PingFederate environment to easily identify it in the upper right corner of the console. If you have deployed PingFederate in a clustered environment, you will also find your nodes on the dashboard with various information at a glance, organized by their respective regions if applicable. Lastly, the new experience comes with a **Search** bar, where you can find and reach your menu item with just a few keystrokes. To see a list of menu items, see *Navigation tabs and menus* on page 19.

Single sign-on to the administrative console

In addition to the new administrative experience, you can now enable single sign-on for the console. You can configure PingFederate itself as the OpenID Provider, in which case you can create an authentication policy for console access. Alternatively, because the implementation is based on the OpenID Connect specifications, you may use any other standards-compliant OpenID Connect Provider.

OAuth 2.0 authorization to the administrative API

The PingFederate administrative API now supports the OAuth 2.0 authorization framework. When enabled, client applications can simply include an access token in each request as opposed to providing authentication information. This modernization makes it easier and more secure for developers who are familiar with OAuth 2.0 to develop applications that update PingFederate configuration.

Authentication API

Mobile applications can now orchestrate authentication via REST APIs without HTTP redirections. Equipped with this redirectless capability, applications have complete control over the authentication experience, which can help improve user-engagement.

Customer IAM

Authentication sessions control when PingFederate reinvokes authentication sources for users that have previously authenticated. Starting with version 10.1, PingFederate is capable of creating authentication sessions upon user registration, which streamlines the overall user experience.

Additionally, the PingFederate authentication API now supports self-service user registration. With this new enhancement, you can offer your customers a seamless experience as they complete their registrations and potentially access multiple applications.

OAuth

Session Management API

The Session Revocation API allows applications to revoke authentication sessions and to check their revocation status. It does not disclose whether a session is valid or any further information. While this use case remains valid, we are introducing a new API that offers more in version 10.1.

The new Session Management API allows applications to leverage the session identifier to query, extend, or revoke authentication sessions of the users. For example, an application can gain knowledge of whether a session is valid and the relevant information, such as idle timeout and session lifetime, if the session is valid. This new capability allows applications to deliver the desired contents and experiences at different times based on the status of sessions. Like the access to the Session Revocation API, the access to the new Session Management API is configurable on a per client basis.

Session identifier in access tokens

Access tokens can now include the session identifier regardless of whether the Access Token Managers are configured to conjoin the validity of access tokens and the authentication sessions of the users. This new capability enables applications to utilize the new Session Management API as well as the preexisting Session Revocation API, as they both require the session identifier. PingFederate administrators can configure this option on a per-Access Token Manager basis.

Centralized signing keys

PingFederate can now use centralized static or dynamically rotating keys to sign self-contained (JWT) access tokens. If enabled, clients can retrieve the signing keys at the standard /pf/JWKS endpoint for the purpose of validating the integrity of the access tokens issued by PingFederate.

Dynamic client registration management protocol

To support organizations and industries that continue to adopt the Dynamic Client Registration Protocol, we are introducing support for the Dynamic Client Registration Management Protocol (*RFC7592*), which defines an open standard for developers to manage their registered clients. When enabled, developers can retrieve and update client information on their own; they can also remove client records if you allow them to do so.

Third-party cryptographic software support

PingFederate 10.1 supports Bouncy Castle FIPS as the provider of its Java keystore and cryptographic operations. In Bouncy Castle FIPS mode, whenever PingFederate uses FIPS-approved algorithms, it uses the Bouncy Castle implementation of those algorithms. There are still a number of cases where PingFederate uses algorithms that are not FIPS-approved. For details on the contexts where PingFederate uses algorithms that are not FIPS-approved, contact customer support.

Enhancements for clustered environments

Improved configuration replication time

When consuming configuration data from the console, engine nodes previously paused requests in a queue and resumed processing of requests once the configuration data was applied. While this wait time was typically very brief, it could be longer in some complex environments, which could have a negative impact on the end-user experience. PingFederate 10.1 reduces this wait time dramatically, eliminating end user impact during configuration replication.

Per-SSO transaction state replication

PingFederate 10.1 offers an optional capability to replicate transaction states of a single request from the region that processes the request to other regions in a clustered environment. It solves the rare but possible problem where the initial authentication request arrives at one data center (defined as an adaptive clustering node group) but subsequent requests are fulfilled at another.

Correlation identifier between PingFederate and PingDirectory

Starting with version 10.1, PingFederate sends the tracking IDs of runtime requests, which is always recorded in the PingFederate audit log, to PingDirectory in LDAP requests. As PingDirectory processes these requests, it records the tracking IDs as sessionID values in its access log. There is also a related capability where request parameters received at certain PingFederate endpoints can be sent to PingDirectory to be recorded as requestID values in its access log.

Because the correlation identifiers are recorded on both sides, reviewing log messages from both products can be useful for auditing and troubleshooting purposes.

Administrative API enhancements

The administrative API has been extended to include the following workflows:

- /incomingProxySettings to manage incoming proxy settings.
- /protocolMetadata to manage meta signing and lifetime settings.
- /serviceAuthentication to manage service authentication settings and credentials.

Other improvements

PingFederate 10.1 introduces a new administrative role: Expression Admin. This role enables
users with the administrative role of Admin to add new or update existing attribute mapping
expressions, expression-driven issuance criteria, and message customization. Without the
Expression Admin role, any Admin-only users are not allowed to introduce new expressions
or to replace the current expression with another one.

- The /idp/startSLO.ping application endpoint now supports a new query parameter,
 LogoutType, to provide a customizable single logout (SLO) experience. This query parameter supports three values: AsyncOnly, SyncOnly, and All.
 - AsyncOnly will cause PingFederate to only log out OpenID Connect and WS-Federation partners in parallel.
 - SyncOnly will cause PingFederate to only log out SAML partners and adapters in series.
 - All will cause PingFederate to log out SAML partners and adapters in series first, followed by OpenID Connect and WS-Federation partners in parallel. Not providing the LogoutType query parameter also has the same effect because this is the default SLO behavior.
- Updated the following bundled components and third-party dependencies:
 - OpenToken Adapter 2.6.2
 - UnboundID LDAP SDK 4.0.13
 - Jackson databind 2.9.10.4
 - Spring Framework 4.3.24
 - JQuery 3.4.1

Ticket ID	Description			
PF-26336	Resolved an issue that could stop PingFederate from responding while it initializes a newly configured plugin.			
PF-26330	Fixed a problem in the administrative console that caused an error when clicking the Manage Certificates button on the Allowed Issuer Certificates window under WS-Trust Mutual SSL Authentication.			
PF-26073	Changes to the <pf_install>/pingfederate/server/default/data/config-store/com.pingidentity.crypto.jwk.MasterKeySet.xml file are now preserved during an upgrade.</pf_install>			
PF-26004	HTTP tracked parameters now work correctly with WS-Federation SP connections.			
PF-25840	The HTML Form Adapter login template now disables the Sign On button while signing on to prevent errors caused by repeated attempts.			
PF-25791	The JIT provisioning configuration no longer gets lost when saving IdP connection changes.			
PF-25665	The HTML Form Adapter validation now ensures that password change and reset options can only be enabled if at least one underlying Password Credential Validator instance configuration supports them.			
PF-25540	All cookies set by the PingFederate administrative console now have the Secure flag.			
PF-25498	When an invalid assertion during client assertion authentication occurs, the HTTP status code 400 is now correctly returned.			
PF-25392	A problem causing the Fail if Policy Engine Finds No Authentication Source option not to display when you do not have the Service Provider role enabled has been resolved.			
PF-25367	When using PingOne for Enterprise with PingFederate Bridge, reconfiguring the PingOne connection is no longer necessary to enable provisioning at a later time if you did not enable it during setup.			

Ticket ID	Description		
PF-25313	An error in the <pf_install>/ pingfederate/server/default/conf/ template/mail-notifications/message-template-username- recovery.html template has been corrected. The greetUser and successMessage variables are now formatted correctly.</pf_install>		
PF-25259	When configuring an SMTP server in the Notification Publisher, any hostname containing valid characters is now accepted, including ones using a non-public suffix.		
PF-25207	Fixed a problem causing the OAuth client to fail if an OAuth client Extended property length was greater than 2048 bytes.		
PF-25205	Fixed a problem preventing some OAuth grant datastores from updating after an archive import.		
PF-25083	PingFederate no longer issues a warning in the server.log when validating access tokens that are issued through the Client Credentials flow through the introspection endpoint.		
PF-24477	Fixed a problem during JSON-token management configuration that allowed the addition of a certificate even when the user-defined Key ID had been used previously.		
PF-24404	Improved the error message returned when an invalid attribute is used while creating or editing an SP connection through the administrative API.		
PF-22573	Messages have been added to the server log when the IRSM state map is purged because it filled up too quickly.		
PF-15649	The outbound provisioning channel timeout now behaves more consistently.		

Navigation tabs and menus

PingFederate 10.1 features new navigation tabs and menus. When you select a tab at the top of the window, the relevant menus appear on the left menu pane. When you select a menu, the menu items appear. The names of some menu items and windows have changed so they match.

Each navigation tab provides access to multiple menus on the left menu pane

Authentication tab menus	Applications tab menus	Security tab menus	System tab menus
IntegrationPoliciesOAuthToken Exchange	IntegrationOAuthToken Exchange	 Certificate & Key Management System Integration 	 Data & Credential Stores Server OAuth Settings External Systems Monitoring & Notifications Protocol Metadata

Navigation tabs, menus, and menu items in alphabetical order

Tab	Menu	Menu item and window name
Applications	Integration	Adapter-to-Adapter Mappings
		Policy Contract Adapter Mappings

Tab	Menu	Menu item and window name
		SP Adapters
		SP Connections
		SP Default URLs
		Target URL Mapping
	OAuth	Access Token Management
		Access Token Mappings
		CIBA Request Policies
		Clients
		OpenID Connect Policy Management
	Token Exchange	Generator Groups
		Processor Policies
		Token Generator Mappings
		Token Generators
		Token Translator Mappings
Authentication	Integration	Authentication API Applications
		IdP Adapters
		IdP Connections
		IdP Default URL
	OAuth	CIBA Authenticators
		IdP Adapter Grant Mapping
		Policy Contract Grant Mapping
		Resource Owner Credentials Grant Mapping
	Policies	Local Identity Profiles
		Policies
		Policy Contracts
		Selectors
		Sessions
	Token Exchange	STS Request Parameters
		Token Processors
Security	Certificate & Key Management	Certificate Revocation Checking
		OAuth & OpenID Connect Keys
		Partner Metadata URLs
		Signing & Decryption Keys & Certificates
		SSL Client Keys & Certificates

Tab	Menu	Menu item and window name
		SSL Server Certificates
		System Keys
		Trusted CAs
	System Integration	Incoming Proxy Settings
		Redirect Validation
		Service Authentication
System	Data & Credential Stores	Active Directory Domains/Kerberos Realms
		Data Stores
		Identity Store Provisioners
		Password Credential Validators
	External Systems	CAPTCHA Settings
		Connect to PingOne for Enterprise
		Notification Publishers
		SMS Provider Settings
	Monitoring & Notifications	Runtime Notifications
		Runtime Reporting
	OAuth Settings	Authorization Server Settings
		Client Registration Policies
		Client Settings
		Scope Management
	Protocol Metadata	Attribute Requester Mapping
		File Signing
		Metadata Export
		Metadata Settings
		SP Affiliations
	Server	Administrative Accounts
		Cluster Management
		Configuration Archive
		Extended Properties
		General Settings
		License Management

Known issues and limitations

Known issues

Administrative API

/sp/idpConnections

For IdP connections, the administrative API connection support is limited to Browser SSO, WS-Trust STS, and OAuth Assertion Grant connections. As a result, when updating an IdP connection using the administrative API, it is possible to lose inbound provisioning settings previously configured using the administrative console.

/bulk

Only resource types currently supported by the administrative API are included in the exported data. Resources not yet supported include:

- Identity Store Provisioners
- Inbound provisioning settings from IdP connections
- PingOne for Enterprise settings
- SMS Provider settings
- WS-Trust STS settings

Hardware security modules (HSM)

When using PingFederate with an HSM and Oracle Java 8u261 (or later) or Amazon Corretto Java 8u272 (or later), if an HSM certificate is used as an SSL server certificate, browsers might not be able to connect to the server. To resolve this issue, add RSASSA-PSS to the jdk.tls.disabledAlgorithms list in the java.security file under the JAVA HOME directory.

Known limitations

Updating Java 8 to 11

Updating Java version 8 to version 11 results in an error when PingFederate is already installed and running. To work around this issue, uninstall and reinstall the PingFederate Windows service by running the UninstallPingFederateService.bat and InstallPingFederateService.bat files located in cpf_install>/pingfederate/sbin/wrapper.

Administrative console and administrative API

- Previously, the administrative API did not accurately reflect a Persistent Grant Max Lifetime setting of 29 days (or shorter) with the selection of the Grants Do Not Timeout Due To Inactivity option. As a result, if you have configured such OAuth authorization server settings and have generated a bulk export in version 10.0 through 10.0.2, we recommend that you regenerate a new bulk export after upgrading to version 10.0.3 (or a more recent version). The newly exported data does not contain the aforementioned flaw, and you can safely import it to version 10.0.3 (or a more recent version).
- When enabling mutual TLS certificate-based authentication, administrators often configure a list of acceptable client certificate issuers. When an administrator uses a browser to access the console or the administrative API documentation, PingFederate returns to the browser the list of acceptable issuers as part of the TLS handshake. If the browser's client certificate store contains multiple client certificates, the browser often presents to the user only the certificates whose issuer matches one of the acceptable issuers. However, when PingFederate runs in a Java 11 environment, Chrome presents to the administrator all its configured client certificates, regardless of whether the issuer matches one of the acceptable issuers or not.

- Prior to toggling the status of a connection with the administrative API, an administrator must ensure that any expired certificates or no longer available attributes are replaced with valid certificates or attributes; otherwise, the update request fails.
- When creating or updating a child instance of a hierarchical plugin, the administrative API retains objects with an "inherited": false name/value pair (or without such name/value pair altogether), ignores those with a value of true, and returns a 200 HTTP status code. No error messages are returned for the ignored objects.
- Using the browser's navigation mechanisms (for example, the Back button) causes inconsistent behavior in the administrative console. Use the navigation buttons provided at the bottom of windows in the PingFederate console.
- If authenticated to the PingFederate administrative console using certificate authentication, a session that has timed out might not appear to behave as expected. Normally (when using password authentication), when a session has timed out and a user attempts some action in the console, the browser is redirected to the login page, and then back to the administrative console once authentication is complete. Similar behavior applies for certificate authentication, in principle. However, because the browser might automatically resubmit the certificate for authentication, the browser might redirect to the administrative console and not the login page.

Upgrading PingFederate

It is not possible to perform a rolling upgrade from PingFederate 10.1.0 to a later maintenance release of 10.1. Upgraded nodes will hang on start-up until the entire cluster has been upgraded. This issue only affects 10.1.0. It does not affect upgrading from 10.1.1 or later to a subsequent maintenance release.

Hardware security modules

- PingFederate must be deployed with Oracle Server JRE (Java SE Runtime Environment) 8, or Amazon Corretto 8.
- When using PingFederate with an HSM from Gemalto or nCipher, it is not possible to use an elliptic curve (EC) certificate as an SSL server certificate.
- When using PingFederate with an HSM from Gemalto, it is not possible to generate a self-signed EC certificate or use an EC certificate as a signing certificate. Furthermore, if only GCM cipher suites are enabled, any attempt to connect to PingOne for Enterprise will result in a "Cipher suite cannot be null" error condition.
- When using PingFederate with an HSM from Gemalto in FIPS approved mode, it is not possible to generate an RSA key with a key size of 1,024.

SSO and SLO

- When consuming SAML metadata, PingFederate does not report an error when neither the validUntil nor the cacheDuration attribute is included in the metadata. Note that PingFederate does reject expired SAML metadata as indicated by the validUntil attribute value, if it is provided.
- The anchored-certificate trust model cannot be used with the SLO redirect binding because the certificate cannot be included with the logout request.
- If an IdP connection is configured for multiple virtual server IDs, PingFederate will always use the default virtual server ID for IdP Discovery during an SP-initiated SSO event.

Composite Adapter configuration

SLO is not supported when users are authenticated through a Composite Adapter instance that contains another instance of the Composite Adapter.

Self-service password reset

Passwords can be reset for Microsoft Active Directory user accounts without the permission to change password.

OAuth

PingFederate does not support case-sensitive naming convention for OAuth client ID values when client records are stored in a directory server. For example, after creating a client with an ID value of sampleClient, PingFederate does not allow the creation of another client with an ID value of SampleClient.

It is worth noting that while it is possible to create clients using the same ID values with different casings when client records are stored in XML files, a database server, or custom storage (if implemented), we recommend *not* doing so to avoid potential record migration issues.

Customer identity and access management

Some browsers display a date-picker user interface for fields that have been designed for date-specific inputs. Some browsers do not. If one or more date-specific fields are defined on the registration page or the profile management page (or both), end users must enter the dates manually if their browsers do not display a date-picker user interface for those fields.

Provisioning

- LDAP referrals return an error and cause provisioning to fail if the user or group objects are defined at the DC level, and not within an OU or within the Users CN.
- The totalResults value in SCIM responses indicates the number of results returned in the current response, not the total number of estimated results on the LDAP server.

Logging

If a source attribute has been configured for masking in an IdP adapter or IdP connection and the source attribute is mapped to OAuth's persistent grant USER_KEY attribute, then the USER_KEY attribute will not be masked in the server logs. Other persistent grant attributes will be masked.

Database logging

If PingFederate cannot establish a Java Database Connectivity (JDBC) connection at startup, PingFederate will continue to write log messages to the failover log file, despite the failover and resume configuration. When the JDBC connectivity issue is resolved, restart PingFederate. On restart, PingFederate will start writing log messages to the database.

If PingFederate is able to establish a JDBC connection at startup, PingFederate will be able to write log messages to the failover log when it encounters a JDBC connectivity issue *and* resume writing log messages to the database when it re-establishes the JDBC connection.

RADIUS NAS-IP-Address

The RADIUS NAS-IP-Address is only included in Access-Request packets when the pf.bind.engine.address is set with an IPv4 address. IPv6 is not supported.

Configcopy

- When the **configcopy** tool is used to copy all connections, channels, data sources, adapters, or token translators, the overridden properties are applied to all instances. Proceed cautiously when applying overrides for copy-all operations.
- The configcopy tool supports copying only a single reference for each of the following configuration items that are defined for a given connection: adapter, data source, Assertion Consumer Service URL, Single Logout Service URL, and Artifact Resolution Service URL. When multiple items are associated with a given connection, only the first reference to each is copied.
- The configcopy tool does not support creation of configuration data that does not exist in the source. If an override parameter is set for a parameter that does not exist in the source configuration, the behavior of the target system is not guaranteed.
- The configcopy tool, when used for copying plugin configurations (including adapters, token translators, and other data stores), does not currently support overrides of complex data structures, including tables, extended contract attributes, and masked fields.

When the configcopy tool is used to copy connection data, any SOAP Single logout (SLO) endpoints defined in the source are not copied to the target, even if the SOAP SLO endpoint is the only SLO endpoint defined at the source. These must be manually added to the target.

Deprecated features



Tip:

As of PingFederate 10.1, the Upgrade Utility warns the administrator when deprecated components are being used.

Roles and protocols

Starting with PingFederate 10.1, roles and protocols are always enabled and no longer configurable through the administrative console and API.

S3_PING discovery protocol

Starting with PingFederate 10.1, the S3_PING discovery protocol has been deprecated. Customers running on AWS infrastructure should instead use NATIVE_S3_PING.

Red Hat Enterprise Linux install script

Starting with PingFederate 10.0, the Red Hat Enterprise Linux install script is no longer available. To install PingFederate 10.0 for Linux, you must download and extract the product distribution .zip file.

SNMP

Monitoring and reporting through the Simple Network Management Protocol (SNMP) is deprecated and will be removed in a future release.

Email configuration

Starting with version 9.3, PingFederate provides a pluggable architecture to publish notifications in a variety of ways, configurable based on the types of events and users. PingFederate 9.3 includes both SMTP and AWS Simple Notification Service (SNS) publishers out of the box. As a result, the upgrade process now migrates email server settings from the source installation as an SMTP Notification Publisher instance in the new PingFederate installation. See Configuring an SMTP Notification Publisher instance on page 931 for more information.

Oracle Solaris 10

Starting with version 9.2, Oracle Solaris 10 is no longer included in the PingFederate qualification process. For a list of supported operating systems, see System requirements on page 110.

Auto-Connect

The Auto-Connect feature has been discontinued in PingFederate 9.2. If such usage is detected in the source installation during an upgrade, the upgrade tool warns the administrator about it. Administrators should remove Auto-Connect configuration from the source installation and re-run the upgrade tool.

JMX monitoring support for outbound provisioning

JMX monitoring of outbound provisioning is no longer an option starting with version 9.2. If such usage is detected in the source installation during an upgrade, the upgrade tool warns the administrator about it. No further action is required. Outbound provisioning transactions are written to the provisioner-audit.log file in the new installation. See Outbound provisioning audit logging on page 771 for more information.

Logging configuration

The default logging configuration has been optimized since PingFederate 8.2. As a result, the product distribution no longer includes the terse.example.log4j2.xml file in the $product = \frac{pf_install}{pingfederate}$ server/default/conf directory, starting with version 8.4.

See Log4j 2 logging service and configuration on page 762 for more information.

Plain text email notification templates

Starting with version 8.2, PingFederate has switched the format of its email notification from plain text to HTML. The new HTML-based templates (message-template-*.html) are located in the $< pf_install > / pingfederate / server/default/conf/template/mail-notifications directory. As a result, PingFederate no longer maintains the plain text templates (message-template-*.txt).$

To preserve previous modifications, you must migrate custom changes manually. For more information, see *Copying customized files or settings* on page 144 for more information.

SpSessionAuthnAdapterId and SourceResource (query parameters for the /sp/startSLO.ping endpoint)

Support for the previously optional SpSessionAuthnAdapterId and SourceResource have been dropped in favor of the SLO improvements introduced in version 8.2.

BoneCP as the JDBC connection pool library

As of PingFederate 8.0, support for BoneCP as the JDBC connection pool library has been deprecated and replaced with Apache Commons DBCP[™] 2, which requires JDBC 4.1 or later drivers.

Verify the database-driver JAR files, found in the $<\!pf_install>/pingfederate/server/default/lib directory, meet the minimum version requirement. If you are using JDBC drivers of version 4.0 (or earlier), contact your vendors for the latest drivers and replace the older JDBC database-driver JAR files.$

DSA certificate creation

Starting with PingFederate 7.3, it is no longer possible to create DSA key pairs in the certificate management pages of PingFederate. Import of DSA key pairs continues to be supported.

Username Token Translator

As of PingFederate 7.2, the Username Token Translator has been deprecated and replaced with an integrated Username Token Processor. While the integrated Username Token Processor and the deprecated Username Token Translator can be simultaneously deployed, we recommend you migrate to the new token processor.

See *Migrating to the integrated Username Token Processor* on page 153 for more information.

LDAP Adapter

Starting with PingFederate 7.2, the LDAP Adapter is no longer supported. This adapter was deprecated in PingFederate 6.6 and replaced by the LDAP Username Password Credential Validator (PCV), which can be used with the HTML Form or HTTP Basic Adapters. See *Migrating to the integrated LDAP Username PCV* on page 152 for more information.

Previous releases

For information about enhancements and issues resolved in previous major and minor releases of PingFederate, follow the links below to their release notes. From there you can also access the release notes for older maintenance releases.

- 10.0 December 2019
- 9.3 June 2019
- 9.2 December 2018

• 9.0 December 2017 and previous releases

Introduction to PingFederate

PingFederate is an enterprise federation server for user authentication and standards-based single sign-on for employee, partner, and customer identity types.

PingFederate is an enterprise identity bridge. It enables outbound and inbound solutions for single signon (SSO), federated identity management, customer identity and access management, mobile identity security, API security, and social identity integration. Browser-based SSO extends employee, customer, and partner identities across domains without passwords using standard identity protocols such as SAML, WS-Federation, WS-Trust, OAuth and OpenID Connect, and SCIM.

About identity federation and SSO

Federated identity management, or identity federation, enables enterprises to exchange identity information securely across domains, providing browser-based single sign-on (SSO).

Identity federation also integrates access to applications across distinct business units within a single organization. As organizations grow through acquisitions, or when business units maintain separate user repositories and authentication mechanisms across applications, a federated solution to browser-based SSO is desirable.

This cross-domain, identity-management solution provides numerous benefits, ranging from increased end-user satisfaction and enhanced customer relations to reduced cost and greater security and accountability.

For complete information about identity federation and the standards that support it, see *Supported standards* on page 28.

Service providers and identity providers

Identity federation standards identify two operational roles in a single sign-on (SSO) transaction: the identity provider (IdP) and the service provider (SP).

An IdP might be an enterprise that manages accounts for a large number of users who might need secure access to the web-based applications or services of customers, suppliers, and business partners. An SP might be a SaaS provider or a business-process outsourcing (BPO) vendor wanting to simplify client access to its services.



Secure single sign-on

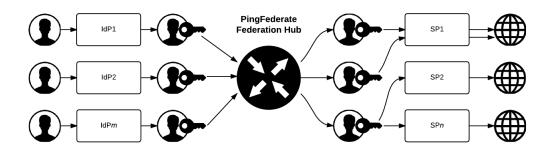
Identity federation allows both types of organizations to define a trust relationship whereby the SP provides access to users from the IdP. The IdP continues to manage its users, and the SP trusts the IdP to authenticate them.

A single instance of PingFederate provides complete support for both roles, even when a single organization's business processes encompass both SP and IdP use cases.

Federation hub

As a federation hub, PingFederate can bridge browser-based single sign-on (SSO) between identity providers and service providers, reducing administrative overhead.

Identity federation refers to the negotiation and management of federation settings with partners. The need to support different federation protocols can hinder application development and SSO implementation. Configuring PingFederate as a federation hub allows you to simplify application development and SSO implementation by extending federated access across partners supporting different federation standards. Configuring PingFederate as a federation hub also provides a centralized console to simplify SSO administration. Bridging identity providers (IdP) and service providers (SP) through a federation hub allows administrators to multiplex a single connection for multiple partners.



Supported standards

PingFederate supports a variety of federation roles, protocols, and standards.

PingFederate provides flexible, integrated support for the SAML protocols, WS-Federation, OAuth, OpenID Connect, and WS-Trust. In addition, PingFederate supports System for Cross-domain Identity Management (SCIM) for inbound and outbound provisioning.

Federation roles

A variety of federation roles work together in an identity federation partnership.

The most recent sets of standards, SAML 2.0 and WS-Federation, define two roles in an identity federation partnership: an identity provider (IdP) and a service provider (SP).



Note:

Earlier SAML 1.x specifications used the terms asserting party (for IdP) and relying party (for SP). For consistency and clarity, PingFederate adopts the later terms IdP and SP across all specifications.

A third role, defined in the SAML 2.0 specifications and available in PingFederate, is that of an IdP Discovery provider.

OAuth 2.0 and OpenID Connect 1.0 can configure PingFederate as an authorization server (AS), an OpenID provider (OP), and a relying party (RP).

Identity provider

An IdP, also called the SAML authority, is a system entity that authenticates a user, or SAML subject, and transmits referential identity information based on the authentication.



Note:

Service provider

An SP is the consumer of identity information provided by the IdP. Based on trust, technical agreements, and verification of adherence to protocols, SP applications and systems determine how to use information contained in an SSO token: a SAML assertion, a JSON Web Token (JWT), or an OAuth access token in conjunction with an ID token.

IdP Discovery provider

This role provides an IdP look-up service that can be incorporated into the implementation of either an IdP or an SP, or employed as a standalone server.

Authorization server

An OAuth authorization server issues access tokens and refresh tokens to OAuth clients after the resource owner fulfills the authentication requirement.

OpenID provider

An OpenID provider (OP) is an AS that is capable of authenticating the resource owner and providing claims (user attributes) to an RP about the authentication event and the user.

Terminology

Definitions for the list of SAML specifications that provide a system of building blocks and support components for achieving secure data exchange in an identity federation.

The list of SAML specifications includes:

- Assertions
- Bindings
- Profiles
- Metadata
- Authentication Context

Assertions

Assertions are XML documents sent from an identity provider (IdP) to a service provider (SP). Each assertion contains identifying information about a user who has initiated a single sign-on (SSO) request.

Bindings

A SAML binding describes the way transport protocols exchange messages. PingFederate supports the following bindings:

HTTP POST

Describes how to transport SAML messages in HTML form-control content, which uses a base-64 format.

HTTP Artifact

Describes how to use an artifact to represent a SAML message. An HTML form control or a query string in the URL transports the artifact.

HTTP Redirect (SAML 2.0)

Describes how to transport SAML messages using HTTP 302 status-code response messages.

SOAP (SAML 2.0)

Describes how to transfer SAML messages across the back channel.

Profiles

Profiles describe processes and message flows combining assertions, request/response message specifications, and bindings to achieve a specific desired functionality or use case. Profiles define the application of the specifications and play a large part in PingFederate.

Metadata

SAML 2.0 defines an XML schema to standardize metadata to facilitate the exchange of configuration information among federation partners, such as profile and binding support, connection endpoints, and certificate information.

Whether you publish or consume metadata, PingFederate supports the use of XML digital signatures to ensure the integrity of the data.

Authentication context

Before allowing access to a protected resource, an SP might want information about the assertion and the user's original authentication by the IdP. The SP uses this information for an access control decision or to provide an audit trail for regulatory or security-policy compliance.

The SAML 2.0 specification provides an XML schema whereby partners create authentication-context declarations. Partners might choose to reference a URI to implement a set of classes provided by the specification to help categorize and simplify context interpretation (see the OASIS document: *Authentication Context for the OASIS Security Assertion Markup Language (SAML) V2.0*). However, it is up to partners to decide if they require additional authentication context and if these classes supply an adequate description. For SAML 1.x, if using the authentication context, called AuthenticationMethod, the context must appear as a URI (see *oasis-sstc-saml-core-1.1.pdf*).

An administrator can configure PingFederate, acting as an IdP, to include a specific authentication context in assertions for browser SSO or WS-Trust.

Several PingFederate integration kits provide methods for the developer to insert authentication context from external IdP applications into the assertion. The SP developer has the option to:

- Call methods for extracting authentication context from an assertion.
- Work with the application to create access control or other processing based on the context.

Browser-based SSO

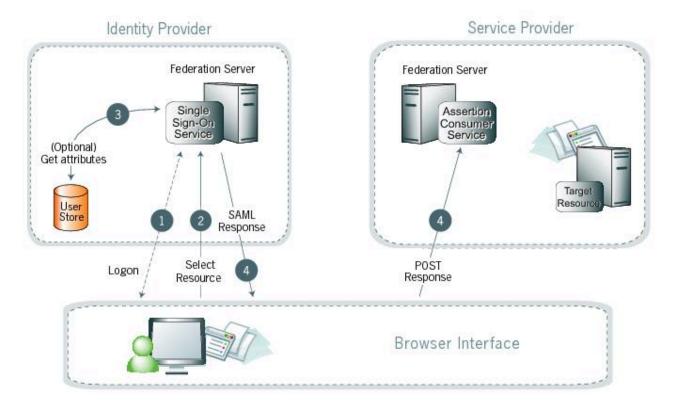
Browser-based single sign-on (SSO) includes *SAML 1.x profiles* on page 30, *SAML 2.0 profiles* on page 33, *WS-Federation* on page 44, and *OpenID Connect* on page 82 and provides standards-based SSO, single logout (SLO), attribute query and X.509 attribute sharing profile (XASP), and the WS-Federation Passive Requestor Profile for service provider (SP)-initiated SSO.

SAML 1.x profiles

SAML 1.0 and 1.1 profiles provide for browser-based single sign-on (SSO), initiated by an identity provider (IdP), using either the POST or artifact bindings.

The SAML 1.x specifications provide for a non-normative service provider (SP)-initiated scenario called "destination-first." This scenario lets web developers create applications that enable a user to initiate SSO from the SP site.

In this scenario, a user logged on to the identity provider (IdP) attempts to access a resource on a remote service provider (SP) server. HTTP POST transports the SAML assertion to the SP.



SSO—Browser-POST profile

Processing steps

- 1. A user logs on to the IdP.
 - If a user is not logged on for some reason, the IdP challenges them to do so at step 2.
- 2. The user clicks a link or otherwise requests access to a protected SP resource.
- 3. Optionally, the IdP retrieves attributes from the user data source.
- 4. The IdP's SSO service returns an HTML form to the browser with a SAML response containing the authentication assertion and any additional attributes. The browser automatically posts the HTML form back to the SP.

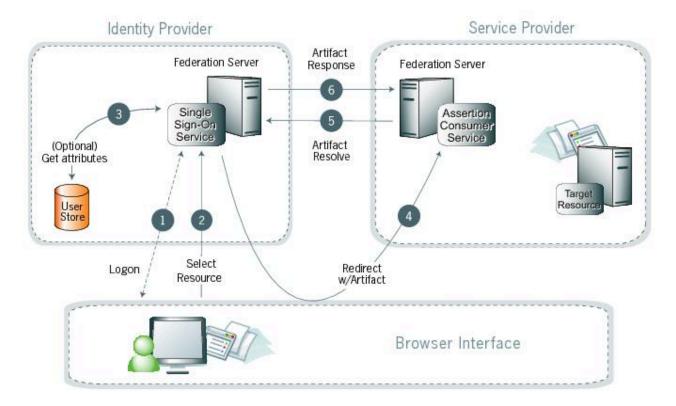


Note:

SAML specifications require digitally-signed POST responses.

5. (Not shown) If the IdP returns a valid SAML assertion to the SP, a session is established on the SP and the browser is redirected to the target resource.

In this scenario, the identity provider (IdP) sends a SAML artifact to the service provider (SP) through either HTTP POST or a redirect. The SP uses the artifact to obtain the associated SAML response from the IdP.



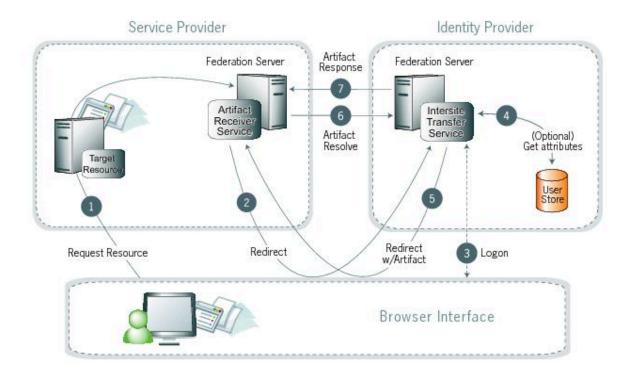
SSO—Browser-Artifact profile

Processing steps

- 1. A user logs on to the IdP.
 - If a user is not logged on for some reason, the IdP challenges them to do so at step 2.
- 2. The user clicks a link or otherwise requests access to a protected SP resource.
- 3. Optionally, the IdP retrieves attributes from the user datastore.
- 4. The IdP federation server generates an assertion, creates an artifact, and sends an HTTP redirect containing the artifact through the browser to the SP's Assertion Consumer Service (ACS).
- 5. The ACS extracts the Source ID from the SAML artifact and sends an artifact-resolve message to the identity federation server's Artifact Resolution Service (ARS).
- 6. The ARS sends a SAML artifact response message containing the previously-generated assertion.
- 7. (Not shown) If the IdP returns a valid SAML assertion to the SP, a session is established on the SP and the browser is redirected to the target resource.

SP-initiated (destination-first) SSO

In service provider (SP)-initiated, destination-first transactions, the user connects to an SP site and attempts to access a protected resource in the SP domain. The user might have an account at the SP site, but according to the federation agreement, the identity provider (IdP) manages authentication. The SP sends an authentication request to the IdP.



SP-initiated SSO

Processing steps

- 1. The user requests access to a protected SP resource. The request redirects to the federation server to handle authentication.
- 2. The federation server sends a SAML request for authentication to the IdP's single sign-on (SSO) service, also called the Intersite Transfer Service.
- 3. If the user is not already logged on to the IdP site or needs to re-authenticate, The IdP asks for credentials, such as ID and password, and the user logs on.
- 4. The user data store can provide additional information about the user for inclusion in the SAML response. The federation agreement between the IdP and the SP predetermines these attributes. See
- 5. The IdP's Intersite Transfer Service returns an artifact representing the SAML response to the SP.
- 6. The SP's artifact handling service sends a SOAP request with the artifact to the IdP's artifact resolver
- 7. The IdP resolves the artifact and returns the corresponding SAML response with the SSO assertion.
- 8. (Not shown) If the IdP returns a valid SAML assertion to the SP, a session is established on the SP and the browser is redirected to the target resource.

SAML 2.0 profiles

PingFederate supports the following major profiles defined under the SAML 2.0 standard:

- Single sign-on on page 34 (SSO)
- Single logout on page 43 (SLO)
- Attribute Query and XASP on page 43
- Standard IdP Discovery on page 44

Single sign-on

Pairing service provider (SP)- and identity provider (IdP)-initiated protocols with transport-binding specifications results in eight practical SSO scenarios.

Enabling SP-initiated transactions through SAML 2.0 increases the number of possible SSO profile variations. The following profile variations each illustrate a specific scenario:

- SP-initiated SSO—POST-POST on page 34 A user attempts to access a protected resource directly on an SP website without logging on. The user does not have an account on the SP site but does have a third-party IdP-federated account. The SP sends an authentication request to the IdP. The user's browser sends both the request and the returned SAML assertion through HTTP POST.
- SP-initiated SSO—Redirect-POST on page 36 The SP sends an HTTP redirect message to the IdP containing an authentication request. The IdP returns a SAML response with an assertion to the SP via HTTP POST
- SP-initiated SSO—Artifact-POST on page 37 The SP sends a SAML artifact to the IdP through an HTTP redirect. The IdP uses the artifact to obtain an authentication request from the SP's SAML artifact resolution service. The IdP returns a SAML response to the SP via HTTP POST.
- SP-initiated SSO—POST-Artifact on page 38 the SP sends an authentication request to the IdP through HTTP POST. The returned SAML assertion is redirected through the user's browser. The response contains a SAML artifact.
- SP-initiated SSO—Redirect-Artifact on page 39 The SP sends an HTTP redirect message to the IdP containing a request for authentication. The IdP returns an artifact through HTTP redirect. The SP uses the artifact to obtain the SAML response.
- SP-initiated SSO—Artifact-Artifact on page 40 The SP sends a SAML artifact to the IdP through an HTTP redirect. The IdP uses the artifact to obtain an authentication request from the SP, and then the IdP sends another artifact to the SP, which the SP uses to obtain the SAML response.
- IdP-initiated SSO—POST on page 41 A user is logged on to the IdP and attempts to access a resource on a remote SP server. The SAML assertion is transported to the SP through HTTP POST.
- IdP-initiated SSO—Artifact on page 42 The IdP sends a SAML artifact to the SP through an HTTP redirect. The SP uses the artifact to obtain the associated SAML response from the IdP.

SP-initiated SSO—POST-POST

In this scenario, a user attempts to access a protected resource directly on a service provider (SP) website without logging on. The user does not have an account on the SP site but does have a third-party identity provider (IdP)-federated account. The SP sends an authentication request to the IdP. The user's browser sends both the request and the returned SAML assertion through HTTP POST.

SP-initiated SSO—POST-POST

Processing steps

- 1. The user requests access to a protected SP resource. The request redirects to the federation server to handle authentication.
- 2. The federation server sends an HTML form back to the browser with a SAML request for authentication from the IdP. The HTML form automatically posts to the IdP's SSO service.
- 3. If the user is not already logged on to the IdP site or needs to re-authenticate, The IdP asks for credentials, such as ID and password, and the user logs on.
- 4. The user data store can provide additional information about the user for inclusion in the SAML response. The federation agreement between the IdP and the SP predetermines these attributes. See
- 5. The IdP's SSO service returns an HTML form to the browser with a SAML response containing the authentication assertion and any additional attributes. The browser automatically posts the HTML form back to the SP.



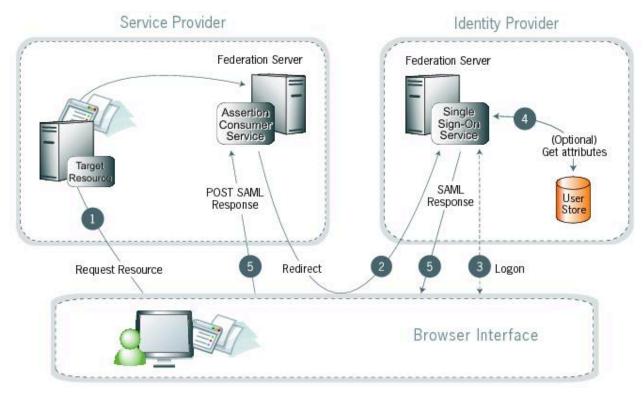
Note:

SAML specifications require digitally-signed POST responses.

6. (Not shown) If the signature and the assertion, or the JSON Web Token, are valid, the SP establishes a session for the user and redirects the browser to the target resource.

SP-initiated SSO—Redirect-POST

In this scenario, the service provider (SP) sends an HTTP redirect message to the identity provider (IdP) containing an authentication request. The IdP returns a SAML response with an assertion to the SP through HTTP POST.



SP-initiated SSO: redirect/POST

Processing steps

- 1. A user requests access to a protected SP resource. The user is not logged on to the site. The request redirects to the federation server to handle authentication.
- 2. The SP returns an HTTP redirect code, 302 or 303, containing a SAML request for authentication through the user's browser to the IdP's single sign-on (SSO) service.
- 3. If the user is not already logged on to the IdP site or needs to re-authenticate, The IdP asks for credentials, such as ID and password, and the user logs on.
- 4. The user data store can provide additional information about the user for inclusion in the SAML response. The federation agreement between the IdP and the SP predetermines these attributes. See
- 5. The IdP's SSO service returns an HTML form to the browser with a SAML response containing the authentication assertion and any additional attributes. The browser automatically posts the HTML form back to the SP.

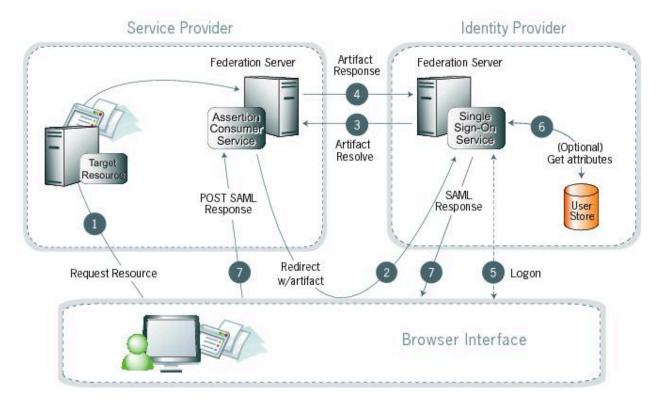


Note:

SAML specifications require digitally-signed POST responses.

6. (Not shown) If the signature and the assertion, or the JSON Web Token, are valid, the SP establishes a session for the user and redirects the browser to the target resource.

In this scenario, the service provider (SP) sends a SAML artifact to the identity provider (IdP) through an HTTP redirect. The IdP uses the artifact to obtain an authentication request from the SP's SAML artifact resolution service (ARS), and the IdP returns a SAML response to the SP through HTTP POST.



SP-initiated SSO—Artifact-POST

Processing steps

- 1. A user requests access to a protected SP resource. The user is not logged on to the site. The request redirects to the federation server to handle authentication.
- 2. The SP generates an authentication request and creates an artifact. The SP sends an HTTP redirect containing the artifact through the user's browser to the IdP's single sign-on (SSO) service.



Note:

The artifact contains the source ID of the SP's artifact resolution service and a reference to the authentication.

3. The SSO service extracts a source ID from the SAML artifact and sends a SAML artifact-resolve message over SOAP containing the artifact to the SP's ARS.



Note:

The federation agreement made prior to this action maps the SP and IdP's source IDs and remote ARS.

- 4. The SP's ARS returns a SAML message containing the previously-generated authentication request.
- 5. If the user is not already logged on to the IdP site or needs to re-authenticate, The IdP asks for credentials, such as ID and password, and the user logs on.

- 6. The user data store can provide additional information about the user for inclusion in the SAML response. The federation agreement between the IdP and the SP predetermines these attributes. See
- 7. The IdP's SSO service returns an HTML form to the browser with a SAML response containing the authentication assertion and any additional attributes. The browser automatically posts the HTML form back to the SP.



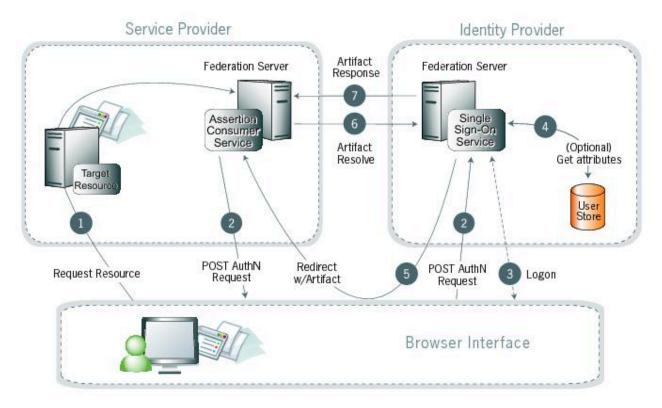
Note:

SAML specifications require digitally-signed POST responses.

8. (Not shown) If the signature and the assertion, or the JSON Web Token, are valid, the SP establishes a session for the user and redirects the browser to the target resource.

SP-initiated SSO—POST-Artifact

In this single sign-on (SSO) scenario, the service provider (SP) sends an authentication request to the identity provider (IdP) through HTTP POST. The returned SAML assertion redirects through the user's browser, and the response contains a SAML artifact.



SP-initiated SSO—POST-Artifact

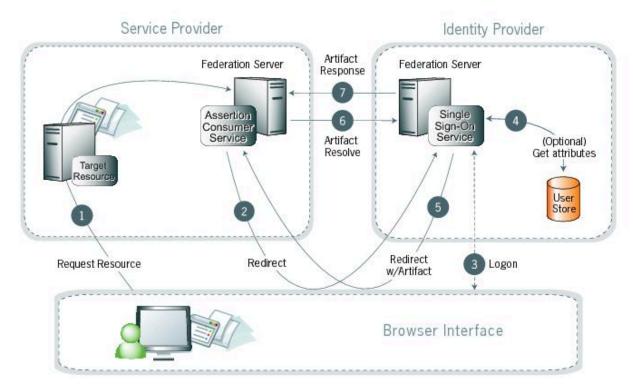
Processing steps

- 1. A user requests access to a protected SP resource. The user is not logged on to the site. The request redirects to the federation server to handle authentication.
- 2. The federation server sends an HTML form back to the browser with a SAML request for authentication from the IdP. The HTML form automatically posts to the IdP's SSO service.
- 3. If the user is not already logged on to the IdP site or needs to re-authenticate, The IdP asks for credentials, such as ID and password, and the user logs on.

- 5. The IdP federation server generates an assertion, creates an artifact, and sends an HTTP redirect containing the artifact through the browser to the SP's Assertion Consumer Service (ACS).
- 6. The ACS extracts the source ID from the SAML artifact and sends an artifact-resolve message to the federation server's Artifact Resolution Service (ARS).
- 7. The ARS sends a SAML artifact response message containing the previously-generated assertion.
- 8. (Not shown) If the IdP returns a valid SAML assertion to the SP, a session is established on the SP and the browser is redirected to the target resource.

SP-initiated SSO—Redirect-Artifact

In this scenario, the service provider (SP) sends an HTTP redirect message to the identity provider (IdP) containing a request for authentication. The IdP returns an artifact through HTTP redirect, and the SP uses the artifact to obtain the SAML response.



SP-initiated SSO—Redirect-Artifact

Processing steps

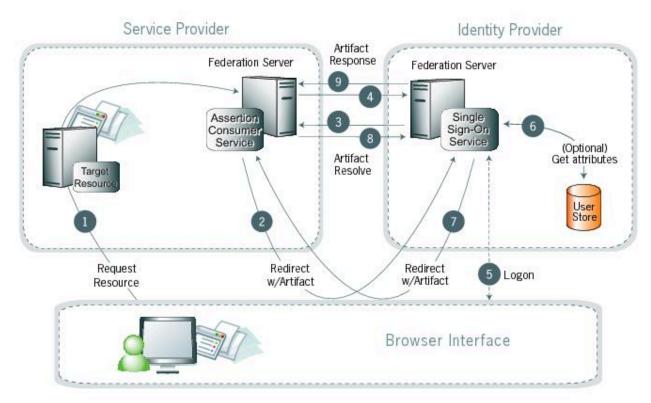
- 1. A user requests access to a protected SP resource. The user is not logged on to the site. The request redirects to the federation server to handle authentication.
- 2. The SP returns an HTTP redirect, either code 302 or 303, containing a SAML request for authentication through the user's browser to the IdP's single sign-on (SSO) service.
- 3. If the user is not already logged on to the IdP site or needs to re-authenticate, The IdP asks for credentials, such as ID and password, and the user logs on.
- 4. The user data store can provide additional information about the user for inclusion in the SAML response. The federation agreement between the IdP and the SP predetermines these attributes. See

.

- 5. The IdP federation server generates an assertion, creates an artifact, and sends an HTTP redirect containing the artifact through the browser to the SP's Assertion Consumer Service (ACS).
- 6. The ACS extracts the Source ID from the SAML artifact and sends an artifact-resolve message to the identity federation server's Artifact Resolution Service (ARS).
- 7. The ARS sends a SAML artifact response message containing the previously-generated assertion.
- 8. (Not shown) If the IdP returns a valid SAML assertion to the SP, a session is established on the SP and the browser is redirected to the target resource.

SP-initiated SSO—Artifact-Artifact

In this scenario, the service provider (SP) sends a SAML artifact to the identity provider (IdP) through an HTTP redirect. The IdP uses the artifact to obtain an authentication request from the SP, then the IdP sends another artifact to the SP, which the SP uses to obtain the SAML response.



SP-initiated SSO—Artifact-Artifact

Processing steps

- 1. A user requests access to a protected SP resource. The user is not logged on to the site. The request redirects to the federation server to handle authentication.
- 2. The Assertion Consumer Service (ACS) generates an authentication request and creates an artifact. It sends an HTTP redirect containing the artifact through the user's browser to the IdP's single sign-on (SSO) service.



Note:

The artifact contains the source ID of the SP's Artifact Resolution Service (ARS) and a reference to the authentication request.

3. The SSO service extracts the source ID from the SAML artifact and sends a SAML artifact resolve message containing the artifact to the SP's artifact resolution service.



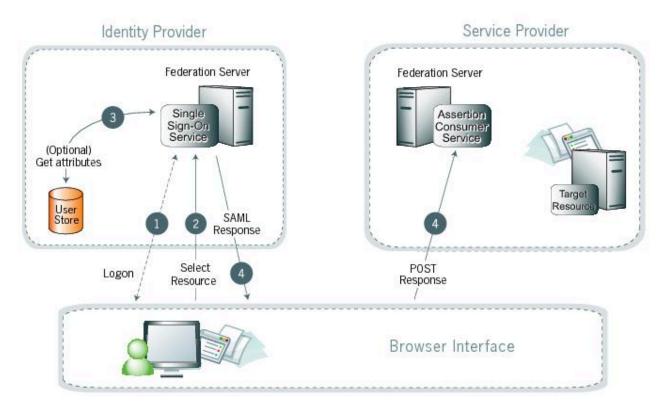
Note:

The federation agreement maps the SP and IdP's source IDs and remote artifact resolution services prior to this action.

- 4. The SP's artifact resolution service sends back a SAML artifact response message containing the previously-generated authentication request.
- 5. If the user is not already logged on to the IdP site or needs to re-authenticate, The IdP asks for credentials, such as ID and password, and the user logs on.
- 6. The user data store can provide additional information about the user for inclusion in the SAML response. The federation agreement between the IdP and the SP predetermines these attributes. See
- 7. The IdP federation server generates an assertion, creates an artifact, and sends an HTTP redirect containing the artifact through the browser to the SP's Assertion Consumer Service (ACS).
- 8. The ACS extracts the Source ID from the SAML artifact and sends an artifact-resolve message to the identity federation server's Artifact Resolution Service (ARS).
- 9. The ARS sends a SAML artifact response message containing the previously-generated assertion.
- 10. (Not shown) If the IdP returns a valid SAML assertion to the SP, a session is established on the SP and the browser is redirected to the target resource.

IdP-initiated SSO—POST

In this scenario, a user is logged on to the identity provider (IdP) and attempts to access a resource on a remote service provider (SP) server. HTTP POST transports the SAML assertion to the SP.



IdP-initiated SSO--POST

Processing steps

- 1. A user logs on to the IdP.
 - If a user is not yet logged on for some reason, he or she is challenged to do so at step 2.
- 2. The user requests access to a protected SP resource.
- 3. After the user requests access, the IdP might also retrieve attributes from the user datastore..
- 4. The IdP's SSO service returns an HTML form to the browser with a SAML response containing the authentication assertion and any additional attributes. The browser automatically posts the HTML form back to the SP.



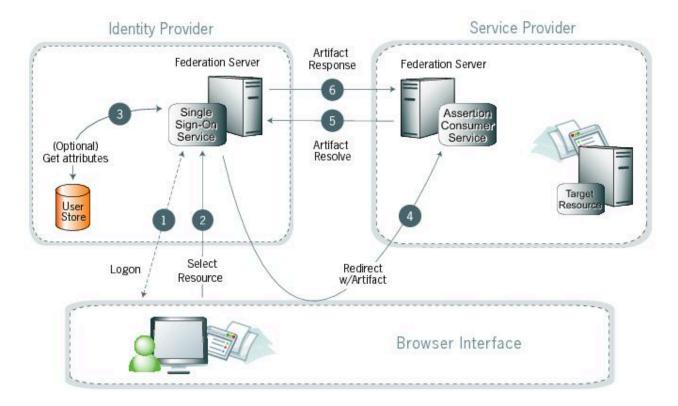
Note:

SAML specifications require digitally-signed POST responses.

5. (Not shown) If the signature and the assertion, or the JSON Web Token, are valid, the SP establishes a session for the user and redirects the browser to the target resource.

IdP-initiated SSO—Artifact

In this single sign-on (SSO) scenario, the identity provider (IdP) sends a SAML artifact to the service provider (SP) through an HTTP redirect. The SP uses the artifact to obtain the associated SAML response from the IdP.



IdP-initiated SSO--Artifact

Processing steps

- 1. A user logs on to the IdP.
 - If a user has not yet logged on for some reason, he or she is challenged to do so at step 2.
- 2. The user clicks a link or otherwise requests access to a protected SP resource.
- 3. After the user requests access, the IdP might also retrieve attributes from the user datastore.

- 4. The IdP federation server generates an assertion, creates an artifact, and sends an HTTP redirect containing the artifact through the browser to the SP's Assertion Consumer Service (ACS).
- 5. The ACS extracts the Source ID from the SAML artifact and sends an artifact-resolve message to the identity federation server's Artifact Resolution Service (ARS).
- 6. The ARS sends a SAML artifact response message containing the previously-generated assertion.
- 7. (Not shown) If the IdP returns a valid SAML assertion to the SP, a session is established on the SP and the browser is redirected to the target resource.

Single logout

The single logout (SLO) profile enables users to log out of all participating sites in a federated session from any site.

The associated identity provider (IdP) federation deployment manages all logout requests and responses for participating sites. If a participating site returns an error, other participating sites might not receive their logout requests. In this scenario, PingFederate returns an error message to the end users.

The logout messages can be transported using any combination of bindings described for SSO (POST, artifact, or redirect). See the diagrams under SAML 2.0 profiles on page 33 for illustrations of these message flows.



When configuring a local loopback connection, in which one PingFederate instance is both the identity provider and the service provider, disable the IdP-Initiated SLO and SP-Initiated SLO options on the Browser SSO window's SAML Profiles tab. These options determine whether SAML logout requests should be sent to the partner during the SLO flow. Those requests aren't necessary and can cause unexpected behavior when the partner connection exists locally. All local sessions for loopback connections are terminated during the SLO flow without the need to send SAML requests.

About session cleanup

When a service provider (SP) receives an SLO request from an IdP, the session creation adapters must handle any session clean-up involving the local application.

Attribute Query and XASP

The SAML 2.0 Attribute Query profile allows a service provider (SP) to request user attributes from an identity provider (IdP) in a secure transaction separate from single sign-on (SSO). The X.509 Attribute Sharing Profile (XASP) defines a specialized extension of the general Attribute Query profile.

The IdP, acting as an attribute authority, accepts attribute queries, performs a datastore lookup into a user repository such as an LDAP directory, provides values to the requested attributes, and generates an attribute response back to the originating SP requester. The SP then returns the attributes to the requesting application.



When privacy is required for sensitive attributes, you can configure PingFederate to obfuscate, or mask, their values in the server and transaction logs.

Web SSO is distinct from the Attribute Query use case. You can configure PingFederate servers to implement either of these profiles without regard to the other.

The XASP specification enables organizations with an investment in Public Key Infrastructure (PKI) to issue and receive Attribute Queries based on user-certificate authentication.

Under XASP a user authenticates directly with an SP application by providing their X.509 certificate. Once the user is authenticated, the SP application requests additional user attributes by contacting the SP PingFederate server. A portion of the user's X.509 certificate is included in the request and can be used

Because the user arrives at the SP server already authenticated, no PingFederate adapter is used in this case.

Standard IdP Discovery

SAML 2.0 identity provider (IdP) Discovery provides a cookie-based look-up, when the IdP is not otherwise specified, to dynamically identify a user's IdP during a service provider (SP)-initiated single sign-on (SSO) event.

This mechanism is helpful in cases where an SP is a hub for several IdPs in an identity federation.



In addition to supporting standard IdP Discovery, PingFederate provides a cross-protocol, proprietary mechanism that allows an SP server to write a persistent browser cookie. The cookie contains a reference to the previous IdP-authentication partner for SSO. For more information, see Configuring IdP discovery using a persistent cookie on page 751.

In the standard scenario, when a user requests access to a protected resource on the SP, commondomain browser cookies are used to determine where a user has previously authenticated. Using this information, a PingFederate server can determine which IdP connection to use for sending an authentication request.

PingFederate can serve in up to three different IdP Discovery provider roles: common domain server, common domain cookie writer, and common domain cookie reader. Each of these roles is necessary to support IdP Discovery. The roles can be distributed across multiple servers at different sites.

Common domain server

In this role, the PingFederate server hosts a domain that its federation partners share in common. The common domain server allows partners to manipulate browser cookies that exist within that common domain. PingFederate can serve in this role exclusively or as part of either an IdP or an SP federation role, or both.

Common domain cookie writer

When PingFederate is acting in an IdP role and authenticates a user, it can write an entry in the common domain cookie, including its federation entity ID. An SP can look up this information on the common domain (not the same location as the common domain server described above).

Common domain cookie reader

When PingFederate is acting as an SP and needs to determine past IdP authentications, it reads the common domain cookie. Based on the information contained in the cookie, PingFederate can then initiate an SSO authentication request using the correct IdP connection.

WS-Federation

PingFederate supports the WS-Federation Passive Requestor Profile for service provider (SP)-initiated single sign-on (SSO), enabling interoperability with Microsoft's Active Directory Federation Service (ADFS).

This profile allows for straightforward redirects and HTTP GET and POST methods to transport SAML assertions or JSON web tokens (JWTs) as security tokens for SSO and logout request and response messages for single logout (SLO).

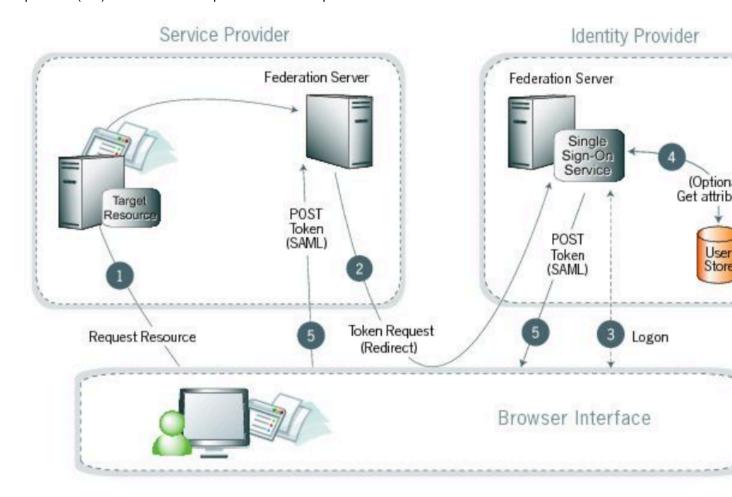


Note:

For more information about WS-Federation and the Passive Requestor Profile, see *web services Federation Languages*.

Passive Requestor profile

This profile permits a user's browser, the passive requestor, to request a security token from an identity provider (IdP) when the user requests access to a protected web service or other resource at an SP.



WS-Federation SSO

Processing steps:

- 1. A user requests access to a protected SP resource. The user is not logged on to the site. The request is redirected to the federation server to handle authentication.
- 2. The SP generates a security token request and redirects the browser to the identity provider's WS-Federation implementation.
- 3. If the user is not already logged on to the IdP site or if re-authentication is required, the IdP asks for credentials, such as ID and password, and the user logs on.
- 4. Additional information about the user can be retrieved from the user datastore for inclusion in the SAML response. These attributes are predetermined as part of the federation agreement between the IdP and the SP. For more information, see *User attributes* on page 92.

- 5. The federation server creates a response containing a signed SAML assertion, or a JSON Web Token, and returns it to the SP through POST.
- 6. (Not shown) If the signature and the assertion, or the JWT, are valid, the SP establishes a session for the user and redirects the browser to the target resource.

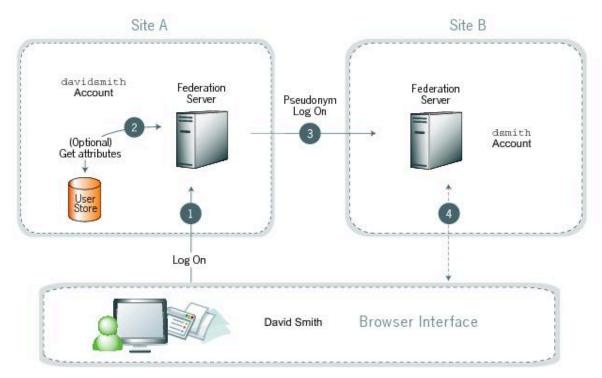
Single logout using WS-Federation is handled in much the same way as with SAML. For more information, see Single logout on page 43. However, HTTP GET/POST is always used as the transport mechanism.

About account linking

Account linking provides a means for a user to log on to disparate sites with just one authentication when the user has established accounts and credentials at each site.

All protocols support this method of interconnecting accounts across domains.

Account linking involves a persistent name identifier associated with accounts at each participating site. The assertion conveys the name identifier, which can be an opaque pseudonym. Once established locally, the service provider (SP) can use the account link to look up the user and provide access without reauthentication.



Account linking

Processing steps

- 1. David Smith logs on to Site A as davidsmith. He then decides to access his account on Site B through Site A.
- 2. Optionally, the federation server looks up additional attributes from the datastore.
- 3. The Site A federation server sends a persistent name identifier to Site B, along with any other attributes.



Note:

When using a pseudonym and sending other attributes, be careful not to send attributes that could identify the subject.

Once the link has been established, Site B stores the information so that David only has to log on to Site A to access Site B.

Web services standards

The PingFederate WS-Trust security token service (STS) interoperates with many different web-service environments that support varying standards.

PingFederate supports multiple versions of SOAP and WS-Trust specifications and operates with any combination of these standards simultaneously.

PingFederate supports namespace aliasing to eliminate common trailing-slash inconsistencies for WS-Trust 1.3. The server does not support namespace aliasing for WS-Trust 2005.

The following table lists supported SOAP/WS-Trust versions and corresponding namespaces.

SOAP/WS-Trust versions

Spec	Version	Namespace
SOAP	1.1 http://schemas.xmlsoap.org/soap/envelope/	
	1.2	http://www.w3.org/2003/05/soap-envelope
WS-Trust 2005 http://schemas.xmlsoa		http://schemas.xmlsoap.org/ws/2005/02/trust/
	1.3	http://docs.oasis-open.org/ws-sx/ws-trust/200512/

Web Services Security

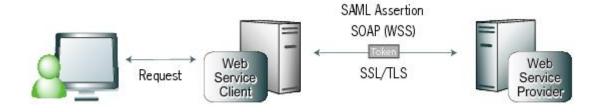
Web Services Security (WSS or WSSE) is a set of specifications defined by the OASIS Web Services Security (WSS) Technical Committee..

WSS defines XML extensions used to secure web service invocations, providing a standard way for partners to add message integrity and confidentiality to web service interactions. The WSS-defined token profiles describe standard ways of binding security tokens to these messages, enabling a variety of additional capabilities. Defined profiles include SAML assertions, Username, Kerberos, X.509, and other existing security tokens. SSL/TLS is often used in conjunction with deployments of WSS. For more information see https://www.oasis-open.org/committees/tc_home.php?wg_abbrev=wss.



Note:

The implementation of WSS in the deployment of web services identity federations is outside the scope of PingFederate, which provides a standalone, standard means of handling the tokens needed for such federations. See *WS-Trust* on page 48.



WSS token transfer

- 1. A user requests content from an application.
- 2. The web service client sends a web service request to the WSP, including the SAML assertion in a WSS header.
- 3. The WSP responds to the request and sends an SSL/TLS token back to the application.
- 4. The web service client returns an HTML page to the user.

WS-Trust

WS-Trust comprises a protocol that systems and applications use when requesting a service to issue, validate, and exchange security tokens.

Organizations can leverage WS-Trust to centralize their security-token processing. The WS-Trust specification defines the role of a Security Token Service (STS) as the entity responsible for responding to requests using the protocol. In this role, the STS creates new security tokens, validates existing security tokens, exchanges security tokens of one type for those of another, or any combination of these.

The WS-SX (Secure Exchange) technical committee manages the WS-Trust specification through its contribution to the OASIS standards organization. For more information, see www.oasis-open.org/committees/tc_home.php?wg_abbrev=ws-sx.

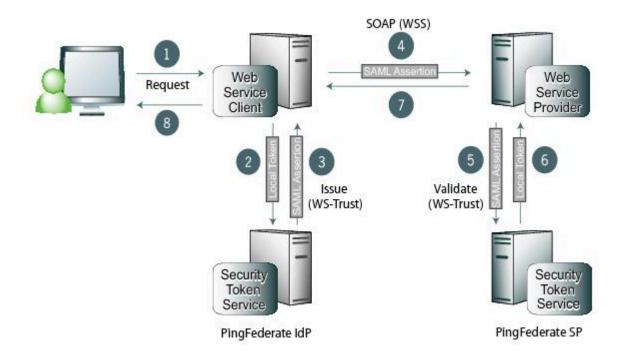
Request types

The WS-Trust protocol defines two request types in securing web services: Issue and Validate, often associated with the web service client (WSC) and web service provider (WSP), respectively.

- The WSC requests that a security token service (STS) issue a SAML token to convey information between the WSC and the WSP.
- The WSP sends the STS a request to validate the incoming token. Optionally, the WSP can request that the STS issue a local token for the service provider (SP) domain.

When issuing and validating security tokens, PingFederate enforces security policies, defined by administrators, generating the token types that are required for a web service request to pass between two security domains (whether these domains are within the same organization or in separate organizations).

Token exchange (example)



- 1. A user requests content from an application.
- 2. The application acts as a WSC to respond to the user's request. The application calls PingFederate, passing the existing user security token to exchange it for the appropriate SAML assertion.
- 3. PingFederate verifies the existing security token, creates a new SAML assertion representing the user, and returns it to the requesting application.
- 4. The application sends a web service request to the WSP, including the SAML assertion in a WS-Security header.
- 5. The WSP retrieves the SAML assertion from the WS-Security header in the incoming request and sends a message to its own deployment of PingFederate to determine if the assertion is valid.
- 6. PingFederate validates the SAML assertion, creates a new security token for the local domain, and returns the new token to the WSP.
- 7. The WSP responds to the request according to its policy for the user.
- 8. The web application returns an HTML page to the user.



Note:

This example shows PingFederate deployed in both the WSC and WSP sides of the interaction. However, other deployment options are also supported.

OAuth 2.0 and PingFederate AS

OAuth 2.0 defines a protocol for securing application access to protected resources by issuing access tokens to clients of REST APIs and non-REST APIs.

Instead of the client directly authenticating to the API using credentials, or the credentials of a user, OAuth enables the client to authenticate by presenting a previously-obtained token. The token represents or contains a set of attributes, policies, or both appropriate to the client and the user. Using these tokens is more secure than using passwords directly on the API call. The attributes are used by the API to authenticate the call and authorize access.

Participants

Client

Wants access to a resource protected by a resource server and interacts with an authorization server to obtain access tokens.

Resource server (RS)

Hosts and protects resources and makes them available to authenticated and authorized clients.

Authorization server (AS)

Issues access tokens and refresh tokens to clients on behalf of the resource servers.

Resource owner (RO)

Denies, grants, or revokes authorization to a client requesting access to resources protected by the resource servers.

Tokens

Access token

Allows clients to authenticate to a resource server and claim authorizations for accessing particular resources. Access tokens have specific authorization scope and duration.

Refresh token

Allows clients to obtain a fresh access token without re-obtaining authorization from the resource owner. A refresh token is a long-lived token that a client can trade in to an authorization server to obtain a new short-lived access token with the same attached authorizations as the existing access token.

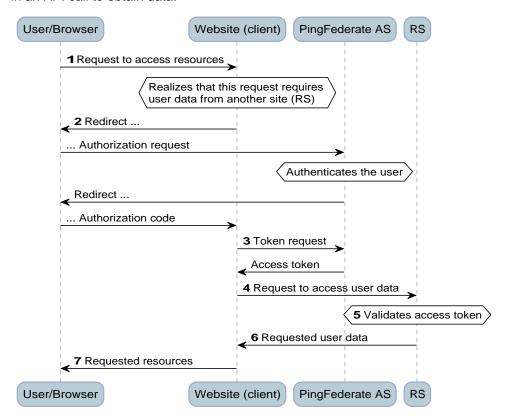
PingFederate OAuth AS

Based on the Internet Engineering Task Force (IETF) OAuth 2.0 Authorization Framework, the OAuth AS in PingFederate supports several interaction models for different types of clients such as servers, desktop applications, or mobile applications. Administrators can also enable Cross-origin Resource Sharing (CORS) support for OAuth endpoints.

Web redirect flow

The web redirect flow process takes place between the user, website, PingFederate authorization server (AS), and resource server (RS).

In this scenario, a user attempts to access a protected resource through a third-party web server client. The client sends an authorization request to the resource server, and receives an authorization code back through a HTTP redirect. The client trades the authorization code for an access token, and uses the token in an API call to obtain data.



Web redirect flow

1. User navigates to an OAuth client website and requests access to protected resources from another website. Flow chart depicting the process of web redirect flow between the User/Browser, Website (client), PingFederate AS and RS.



Note:

To reduce the risk of code interception attach, the OAuth client can optionally include the parameter code challenge with or without code challenge method. For more information, see step 3 and Flow chart depicting the process of web redirect flow between the User/Browser, Proof Key for Code Exchange by (PKCE) OAuth Public Clients.

2. The browser is redirected to the PingFederate OAuth AS with a request for authorization.

If the user is not logged on, the OAuth AS challenges the user to authenticate. The OAuth AS authenticates the user and prompts for authorization. After the user authorizes, the OAuth AS redirects the browser to the requesting site with an authorization code. If the user does not authenticate, the OAuth AS returns an error rather than the authorization code.

3. The requesting site makes a HTTPS request to the OAuth AS to exchange the authorization code for an access token.



If the OAuth client has provided the optional parameter code challenge in step 1, it must submit the corresponding code verifier in this request.

The OAuth AS validates the grant and user data associated with the code and then returns an access

- 4. The requesting site uses the access token in an API call to request user data.
- 5. The RS asks PingFederate for verification that the token is valid and has not expired. PingFederate returns data about the user, the granted scope, and the client ID.
- 6. Once verified, the RS returns the requested data to the requesting site.
- 7. The requesting site displays data from the API call to the user.

Device authorization grant

The device authorization grant process takes place between the user, device, PingFederate authorization server (AS), and resource server (RS).

In this scenario, a user attempts to access a protected resource through a device client that lacks a browser or has limited user-input capabilities, such as a smart TV, digital picture frame, or printer. The OAuth device authorization grant type allows a user to grant authorization to the device client using a browser on a second device, such as a smart phone or computer. For more information about the grant type, see the OAuth 2.0 Device Authorization Grant specification.

OAuth device authorization grant

Processing steps

- 1. The device sends a device authorization request to PingFederate, the authorization server (AS), at its device authorization endpoint.
- 2. PingFederate returns a device authorization response. Among other parameters, the response contains a device code, a user code, a user authorization endpoint, and a user authorization endpoint with the user code in a query parameter.
- 3. The device provides the user authorization endpoint with the user code in a query parameter, the user code, and instructions to the user, as in the following example.

```
Using a browser on another device, visit: https://www.example.com/authorizeDevice

Enter the code: HVF7-B4KW
```

4. The device starts sending device access token requests to PingFederate at its token endpoint to poll whether the user has completed the authorization process.

The device access token request must include the device authorization grant type, urn:ietf:params:oauth:grant-type:device code, the device code, and the user code.

For each device access token request it receives, PingFederate returns a device access token response. The payload varies depending on the authorization status.

- 5. The user completes the authorization process by performing the following actions:
 - Goes to the user authorization endpoint on a second device that has a browser, such as a smartphone or a computer.
 - Fulfills the authentication requirements.
 - c. Enters the user code or confirms a pre-populated user code.
 - d. Approves or denies the scope of permissions requested by the device.
- 6. The device continues polling PingFederate for an authorization status.

If the user denies the scope of permissions, PingFederate provides the device with a relevant error message in the device access token response.

- 8. The device provides the access token to the RS to access protected resources.
- 9. The RS validates the access token.
- 10. The RS provides the requested data to the device.

CIBA grant

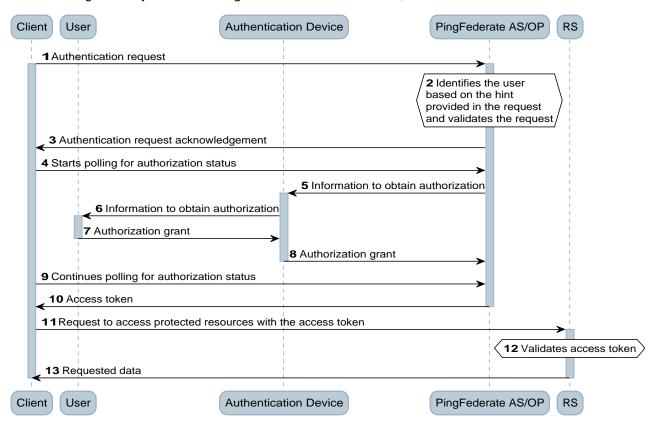
Client Initiated Backchannel Authentication (CIBA) is an extension to OpenID Connect that improves the end-user experience during authentication and authorization in a federated environment.

Like OpenID Connect, CIBA is an authentication flow, governing how clients are identified and granted access. With CIBA, user consent can be requested through an out-of-band flow. For example, when making an online purchase, CIBA improves user experience because the customer's browser will not have to redirect to a financial institution for authorization. Instead, the customer receives a push notification from the financial institution's mobile app to complete authorization. This allows the customer to avoid confusing browser redirects.

CIBA by poll

The OAuth Client Initiated Backchannel Authentication (CIBA) grant by poll process takes place between the client, user, authentication device, PingFederate, and resource server (RS).

After receiving an authentication request acknowledgment, the client starts polling the OpenID Provider (OP)'s token endpoint on a regular interval to obtain the authorization results. When the OP receives the authorization granted by the user through the authentication device, it returns an access token to the client.



OAuth CIBA grant by poll

- 1. The client sends an authentication request to PingFederate at its client-initiated backchannel authentication endpoint. The client must include in its authentication request the desired scope of permissions and one identity hint for PingFederate to identify the user. When providing an identity hint, the client has three options:
 - login hint
 - login_hint_token
 - id token hint

The client can include a user code using the user_code parameter in the authentication request, transmit all request parameters of the authentication request in a signed request object, or do both.

- 2. PingFederate validates the authentication request and identifies the user based on the hint provided by the client.
- 3. PingFederate returns an authentication request acknowledgement to the client. The response contains the identifier, auth_req_i, that PingFederate assigns to the authentication request.
- 4. The client starts polling PingFederate at its token endpoint to check whether the user has completed the authorization process.

The client must include in its token request the CIBA grant type, urn:openid:params:grant-type:ciba, and the corresponding auth_req_id value.

For each token request it receives, PingFederate returns a token response. The payload varies depending on the authorization status.

- 5. PingFederate invokes a CIBA authenticator based on the applicable CIBA request policy to reach out to the user with the information (for example, the requested scopes) that the user needs to obtain authorization.
- 6. The authentication device presents the information and works with the user to obtain authorization.
- 7. The user reviews the information presented by the authentication device and then approves or denies the scopes requested by the client.
- 8. The authentication device sends the authorization result back to PingFederate.
- 9. The client continues polling PingFederate for an authorization result.
- 10. PingFederate returns an access token in a token response to the client.

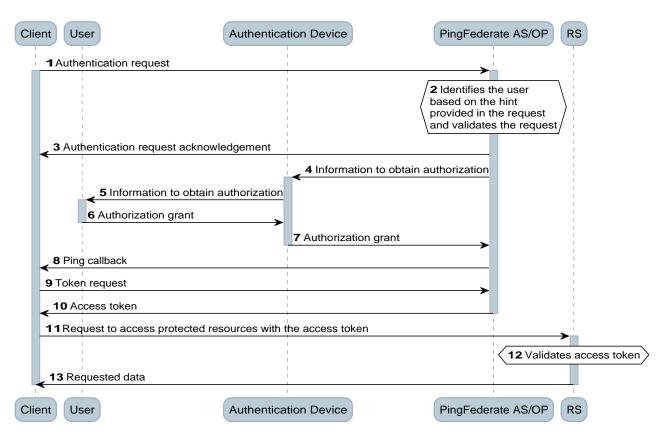
If the user denies the requested scopes, PingFederate provides the client with a relevant error message in the token response.

- 11. The client provides the access token to the RS to access protected resources.
- 12. The RS validates the access token.
- 13. The RS provides the requested data to the client.

CIBA by ping

The OAuth Client Initiated Backchannel Authentication (CIBA) grant by ping process takes place between the client, user, authentication device, PingFederate, and resource server (RS).

After receiving an authentication request acknowledgment, the client waits for a ping callback message from the OpenID Provider (OP). When the OP receives the authorization granted by the user through the authentication device, it sends a ping callback message to the client's notification endpoint. The client then sends a token request to retrieve an access token.



OAuth CIBA grant by ping

Processing steps

1. The client sends an authentication request to PingFederate at its client-initiated backchannel authentication endpoint.

The client must include in its authentication request the desired scope of permissions, one identity hint for PingFederate to identify the user, and a bearer token that PingFederate can use to authenticate the ping callback message. When providing an identity hint, the client has three options:

- login hint
- login hint token
- id token hint

For the bearer token, the client must follow the syntax as defined in RFC 6750, section 2.1 and transmit it using the client notification token parameter.

The client can include a user code using the user code parameter, transmit all request parameters of the authentication request in a signed request object, or do both.

The authentication request can be signed or unsigned.

- 2. PingFederate validates the authentication request and identifies the user based on the hint provided by the client.
- 3. PingFederate returns an authentication request acknowledgment to the client. The response contains the identifier, auth req id, that PingFederate assigns to the authentication request.
- 4. PingFederate invokes a CIBA authenticator based on the applicable CIBA request policy to reach out to the user with the information (for example, the requested scopes) that the user needs to obtain authorization.
- 5. The authentication device presents the information and works with the user to obtain authorization.

- 7. The authentication device sends the authorization result back to PingFederate.
- 8. PingFederate sends a ping callback message using the HTTP POST method to the client at its notification endpoint.

Per specification, PingFederate includes the client_notification_token value in the Authorization HTTP request header and the auth req id value in the message body.

9. The client sends a token request to PingFederate at its token endpoint.

The client must include in its token request the CIBA grant type, urn:openid:params:grant-type:ciba, and the corresponding auth req id value.

10. PingFederate returns an access token in a token response to the client.

If the user denies the requested scopes, PingFederate provides the client with a relevant error message in the token response.

- 11. The client provides the access token to the RS to access protected resources.
- The RS validates the access token.
- 13. The RS provides the requested data to the client.

Token exchange grant

You can configure the PingFederate OAuth server to support the token exchange grant type.

This feature uses the protocol defined in the OAuth 2.0 token exchange specification RFC 8693. For information about the protocol, see *OAuth 2.0 Token exchange*. For information about configuring the OAuth server for token exchange, see *OAuth token exchange* on page 594.

An OAuth token exchange begins when an OAuth client sends a token exchange request to the PingFederate OAuth server. The request contains the token exchange grant type parameter, a subject token, and a subject token type parameter. For impersonation use cases, the request also contains an actor token and actor token type parameter. The request might also contain resource, audience, scope, and requested token type parameters.

The optional resource parameter identifies an Access Token Management (ATM) instance or token generator instance. The optional audience parameter identifies another client that will use the new token to get access to a resource. If a request includes an audience parameter, the OAuth server will select the ATM instance specified in the audience client's configuration as seen in step 2.

A successful token exchange ends when the OAuth server sends the client a message containing the requested token. The response also contains an issued token type parameter and a token type parameter.

When setting up token exchange on the OAuth server, you must create at least one Token Exchange Processor Policy (TEPP). A TEPP includes a map of subject token types to token processor instances and a map of attributes from the request and other sources to a TEPP attribute contract. You must assign a default TEPP to the OAuth server. Optionally, you can assign a TEPP to OAuth clients that need to exchange tokens.

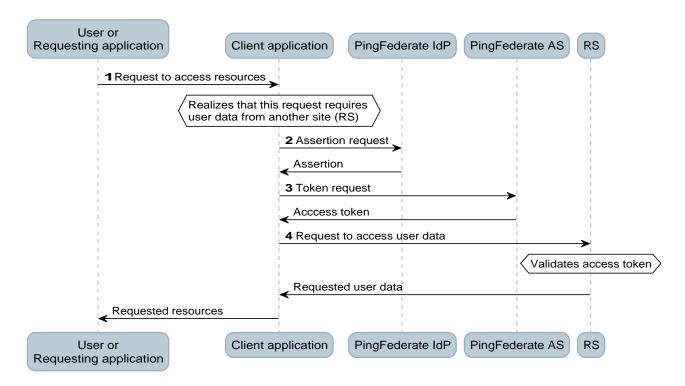
When the OAuth server receives a token exchange request from an OAuth client:

- 1. The OAuth server selects a TEPP and fulfills the TEPP attribute contract as follows:
 - If the OAuth client's configuration specifies a TEPP, the OAuth server selects that TEPP. Otherwise, the OAuth server selects the default OAuth server TEPP.
 - The OAuth server uses the TEPP's map of subject token types to select a token processor instance. The token processor then uses the TEPP's attribute map to produce a TEPP attribute contract.
- 2. The OAuth server selects a token generator instance or ATM instance as follows:
 - If the request does not have an audience or resource parameter, the OAuth server selects the default OAuth server ATM instance.
 - If the request includes resource parameters, audience parameters, or both, and if OAuth server determines that they identify a single ATM or token generator instance, the OAuth server selects it.
 - If the request includes resource parameters, audience parameters, or both, and if OAuth server
 determines that they do not identify a single ATM or token generator instance, the OAuth server
 returns an error.
- 3. The token generator or ATM uses the attribute contract to generate the token as follows:
 - The default token type is generated if the request does not include a requested token type parameter.
 - If the request includes a requested token type parameter, that type of token is generated.
 - If the request includes a requested token type but the token generator or ATM does not support that token type, the OAuth server returns an error.
- 4. The OAuth server sends the client a response containing the requested token.

Assertion grant profile for OAuth 2.0 authorization grants

The assertion grant profile process takes place between the user or requesting application. client application, PingFederate identity provider (IdP), PingFederate authorization server (AS) and resource server (RS).

In this scenario, a client obtains an assertion, either a SAML 2.0 bearer assertion or a JSON Web Token (JWT) bearer token, and makes an HTTP request to the PingFederate OAuth AS to exchange the assertion for an access token. The OAuth AS validates the assertion and returns an access token. The client uses the token in an API call to the RS to obtain data.



Assertion grant profile

Processing steps

- 1. A user-initiated or client-initiated event, such as a mobile application or a scheduled task, requests access to software as a service (SaaS) protected resources from an OAuth client application.
- 2. The client application obtains an assertion from an IdP.



Note:

When using SAML assertions as authorization grants, client applications must obtain assertions that meet the requirements defined in RFC7522. Do not use SAML assertions acquired through browser single sign-on (SSO) profiles here.

- 3. The client application makes an HTTP request to the PingFederate OAuth AS to exchange the assertion for an access token. The OAuth AS validates the assertion and returns the access token.
- 4. The client application adds the access token to its API call to the RS. The RS returns the requested data to the client application.

OpenID Connect support

As an extension of OAuth capabilities, PingFederate supports an optional configuration for OpenID Connect.

OpenID Connect is a modern protocol for secure, lightweight transfer of authentication and user attributes. For more information, see openid.net/connect.

PingFederate can be deployed as an OpenID Provider (OP), a Relying Party (RP), or both. PingFederate supports both the Basic Client and the Implicit Client profiles.

Client management

PingFederate provides features that support OAuth client management.

PingFederate provides administrators the flexibility to manage OAuth clients using the following interfaces:

The administrative console

- The administrative API
- The OAuth Client Management Service

Additionally, PingFederate supports dynamic client registration based on the OAuth 2.0 Dynamic Client Registration Protocol specification.

System for Cross-domain Identity Management (SCIM)

PingFederate supports the SCIM 1.1 protocol for outbound and inbound provisioning.

At an identity provider (IdP) outbound site, you have the option to automatically provision and maintain user accounts at service provider (SP) sites that have implemented SCIM. When you have PingFederate configured as an SP inbound site, you can automatically provision and manage user accounts and groups for your own organization using the standard SCIM protocol. For a brief summary of supported features, see the following table.

Feature	Outbound provisioning	Inbound provisioning
SCIM specification	SCIM 1.1	SCIM 1.1
Data format	JSON	JSON
User and group create, read, update, and delete (CRUD) operations	Yes	Yes
Custom schema support	Yes	Yes
List/query and filtering support	Not applicable	Yes
PATCH	Yes	No
Authentication method	HTTP Basic and OAuth Resource Owner Password Credentials grant type	HTTP Basic and client certificate (mutual TLS)
Source data stores	PingDirectory, Microsoft Active Directory, Oracle Directory Server Enterprise Edition, and Oracle Unified Directory	Not applicable
Target data stores	Not applicable	Active Directory and other data stores via the Identity Store Provisioner Java SDK interface

For detailed information about SCIM, see www.simplecloud.info.

Transport and message security

Two main ways of securing interactions are Secure Sockets Layer with Transport Level Security (SSL/TLS) and digital signatures.

Use SSL/TLS in environments the require both message confidentiality and integrity. For SAML messaging, digital signatures ensure the identity of both parties involved in the transaction and validate that a particular partner received a message. To achieve increased privacy, PingFederate also lets you encrypt SAML 2.0 messages, including SAML metadata files, as well as WS-Trust STS assertions.

For more information, see Security and Privacy Considerations for the OASIS Security Assertion Markup Language (SAML) V2.0.

SSO integration overview

To complete the implementation of a federated-identity network, you must integrate PingFederate programmatically with end-user applications and identity management (IdM) systems. Documentation for integration kits is available on the Ping Identity website.

This content provides an overview of the approaches to integrating systems and applications with PingFederate for browser-based single sign-on (SSO).

To enable both the identity provider (IdP) and service provider (SP) sides of this integration, Ping Identity provides integration kits, including various connectors for secure single sign-on (SSO) to software as a service (SaaS) providers. You can download these integration kits from PingFederate Downloads.

PingFederate also includes an SDK that software developers can use to write custom interfaces for specific systems. For more information, see the SDK Developer's Guide on page 1021, available in the PingFederate distribution sdk directory.

For integration with the PingFederate WS-Trust security token service (STS), we provide a range of token translators. These plugin token processors for an IdP, and token generators for an SP, connect the STS with web service providers and clients for access to identity-enabled web services.

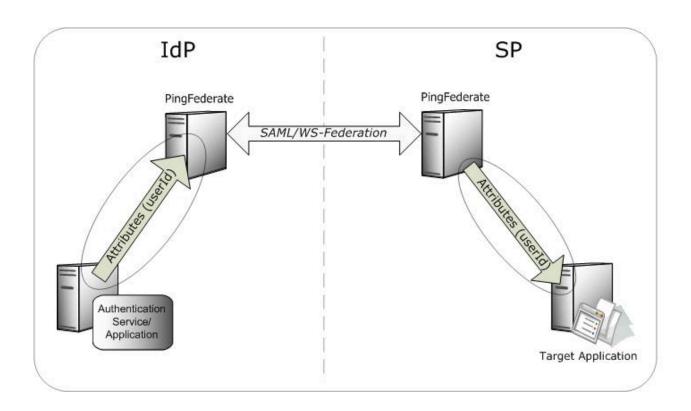
SSO integration concepts

PingFederate supports both identity provider (IdP) and service provider (SP) integration.

For an IdP, the first step in the integration process involves sending identity attributes from an authentication service or application to PingFederate. PingFederate uses those identity attributes to generate a SAML assertion. For information about SAML, see Supported standards on page 28. IdP integration typically provides a mechanism through which PingFederate looks up a user's current authenticated session data, such as a cookie, or authenticate a user without such a session.

For an SP, the last step of the integration process involves sending identity attributes from PingFederate to the target application. PingFederate extracts the identity attributes from the incoming SAML assertion and sends them to the target application to set a valid session cookie or other application-specific security context for the user.

The following diagram illustrates the basic concepts of integration with PingFederate.



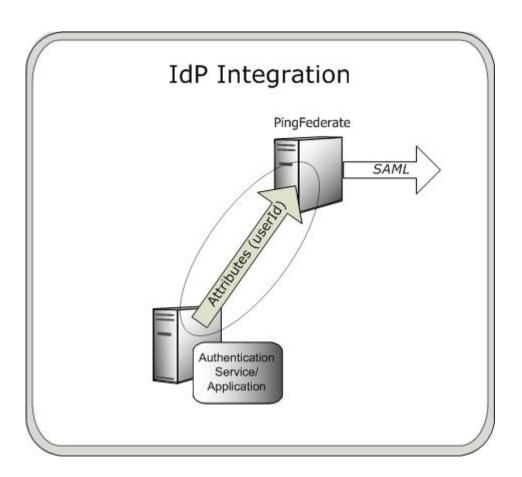
Identity provider integration

Identity provider (IdP) integration involves retrieving user-identity attributes from the IdP domain and sending them to the PingFederate server.

An IdP is a system entity that authenticates a user, or "SAML subject," and transmits referential identity attributes based on that authentication to PingFederate. Typically, the IdP retrieves the identity attributes from an authenticated user session. Depending on the IdP deployment/implementation environment, a number of attribute-retrieval approaches are used for IdP integration. Ping Identity offers a broad range of commercial integration kits that address various IdP scenarios, such as custom-application integration, integration with a commercial identity management (IdM) product, or integration with an authentication system.



For IdPs implementing single sign-on (SSO) to selected software as a service (SaaS) providers such as Google Apps and Salesforce, PingFederate also provides automated user provisioning.



Custom applications

Many applications use their own authentication mechanisms, typically through a database or LDAP repository, and are responsible for their own user-session management. Custom-application integration is necessary when there is limited or no access to the web or application server hosting the application. Integration with these custom applications is handled by application-level integration kits, which allow software developers to integrate their applications with a PingFederate server acting as a service provider (SP).

With these integration kits, PingFederate sends the identity attributes from the SAML assertion to the SP application, which then uses them for its own authentication and session management. As for the IdP, application-specific integration kits include an SP agent, which resides with the SP application and provides a simple programming interface to extract the identity attributes sent from the PingFederate server. PingFederate cam use this information to start a session for the SP application.

Ping Identity provides custom-application integration kits for a variety of programming environments, including:

- Java
- .NET
- PHP

Ping Identity provides an Agentless Integration Kit, which allows developers to use direct HTTP calls to the PingFederate server to temporarily store and retrieve user attributes securely, eliminating the need for an agent interface.

IdM systems

An IdP enterprise that uses an IdM system expands the reach of the IdM domain to external partner applications through integration with PingFederate. IdM integration kits typically use the IdM agent API to access identity attributes in the IdM proprietary session cookie and transmit those attributes to the PingFederate server.

IdM integration kits do not require any development; the PingFederate administrative console accomplishes all integrations.

Ping Identity provides integration kits for many of the leading IdM systems, such as Oracle Access Manager.

Authentication systems

An authentication application or service normally handles initial user authentication outside of the PingFederate server. To access applications outside the security domain, PingFederate authenticationsystem integration kits leverage this local authentication.

Authentication integration kits do not require any development; the PingFederate administrative console accomplishes all integrations with PingFederate. Ping Identity offers integration kits for authentication systems including:

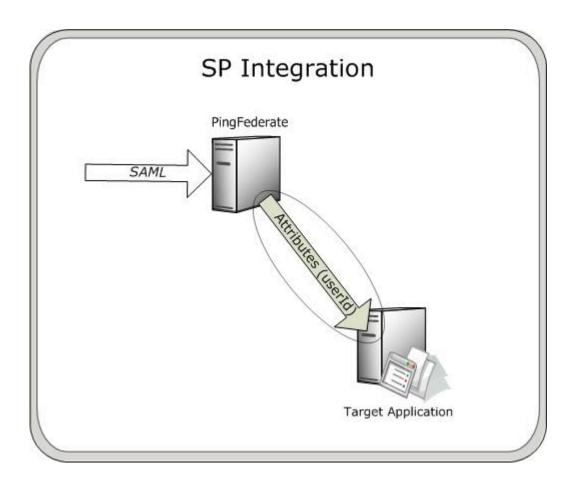
- X.509 Certificate
- RSA SecurID Integration Kit
- Symantec VIP Integration Kit

PingFederate also packages two IdP adapters, an HTML Form Adapter and an HTTP Basic Adapter, which delegate user authentication to plugin password credential validators (PCVs). Supplied validators use either an LDAP directory, RADIUS server, or a simple username/password verification system maintained by PingFederate. Customized validators can also be developed. When the PingFederate IdP server receives an authentication request for SP-initiated SSO or a user clicks a link for IdP-initiated SSO, PingFederate invokes the implemented adapter and prompts the user for credentials, if the user is not already logged on.

Service provider integration

Service provider (SP) integration involves passing the identity attributes from PingFederate to the target SP application.

An SP is the consumer of identity attributes provided by the identity provider (IdP) through a SAML assertion. The SP application uses this information to set a valid session or other security context for the user, represented by the identity attributes. Session creation involves a number of approaches. For the IdP, Ping Identity offers commercial integration kits that address the various SP scenarios. Most SP scenarios involve custom-application integration, server-agent integration, integration with an identity management (IdM) product, or integration with a commercial application.



Custom applications

Many applications use their own authentication mechanisms, typically through a database or LDAP repository, and are responsible for their own user-session management. Custom-application integration is necessary when there is limited or no access to the web or application server hosting the application. Application-level integration kits handle integration with these custom applications and allow software developers to integrate their applications with a PingFederate server acting as an SP.

With these integration kits, PingFederate sends the identity attributes from the SAML assertion to the SP application, which can then use them for its own authentication and session management. As for the IdP, application-specific integration kits include an SP agent, which resides with the SP application and provides a simple programming interface to extract the identity attributes sent from the PingFederate server. PingFederate cam use this information to start a session for the SP application.

Ping Identity provides custom-application integration kits for a variety of programming environments, including:

- Java
- .NET
- PHP

In addition, Ping Identity provides an Agentless Integration Kit, which allows developers to use direct HTTP calls to the PingFederate server to temporarily store and retrieve user attributes securely, eliminating the need for an agent interface.

Server agents

Server-agent integration with PingFederate allows SP enterprises to accept SAML assertions and provide single sign-on (SSO) to all applications running on that web or application server; there is no need to integrate each application. Since integration occurs at the server level, server-agent integration maximizes With server-agent integration kits, PingFederate sends the identity attributes from the SAML assertion to the server agent, which is typically a web filter or Java Authentication and Authorization Service (JAAS) Login Module. The server agent extracts the identity attributes, which the server then uses to authenticate and create a session for the user.

SP server-integration kits do not require any development work: the PingFederate administrative console accomplishes all integrations with PingFederate.

Ping Identity provides integration kits for many web and application servers, including:

- Internet Information Services (IIS)
- Apache (Red Hat)
- Apache (Windows)
- NetWeaver
- WebSphere

IdM systems

IdM integration with PingFederate allows an SP enterprise to accept SAML assertions and provide SSO to applications protected by the IdM domain. IdM integration kits typically use the IdM agent API to create an IdM proprietary session token based on the identity attributes received from PingFederate.

IdM integration kits do not require any development; the PingFederate administrative console and the IdM administration tool accomplish integration with PingFederate.

Ping Identity provides integration kits for leading IdM systems, such as Oracle Access Manager.

Commercial applications and SaaS

Commercial-application integration with PingFederate allows an SP enterprise to accept SAML assertions and provide SSO to those commercial applications.

These integration kits do not require any development; the PingFederate administrative console accomplishes all integrations.

Ping Identity offers integration kits to many commercial applications and SaaS vendors, including:

- Citrix
- SharePoint
- Box
- Google
- Office 365
- Salesforce
- Slack
- Workday
- Zendesk

Integrations and deployment scenarios

Ping Identity provides integrations that support many PingFederate deployment scenarios.

Here are some of the deployment scenarios that the integrations support:

Identity management

Identity management integrations allow PingFederate to extend the domain of identity systems to include the partner applications that you integrate with PingFederate.

Authentication systems

Authentication systems allow users to authenticate with PingFederate through a variety of methods, such as web forms and certificates.

Custom applications

Custom application integrations allow PingFederate to extend single sign-on capabilities to applications that may not have access to a web or application server. They support a variety of programming languages, including Java, .NET, and PHP.

Server agents

Server agents allow PingFederate to extend single sign-on abilities to applications running on a variety of web servers.

Multi-factor authentication (MFA)

MFA integrations allow PingFederate to include third-party MFA providers as part of the sign-on flow.

Mobile device management (MDM)

Mobile device management integrations allow PingFederate to adjust the sign-on flow based on device information.

Risk/Intelligence

Risk intelligence integrations allows PingFederate to retrieve a security risk assessment when a user signs on. You can use this information to dynamically adjust authentication requirements based on the risk level for each sign-on event.

Provisioning

Provisioning connectors allow PingFederate to propagate users and groups from a user directory to a SaaS service. Connectors also include single sign-on integration with the service.

Social login

Cloud identity connectors allow PingFederate to use third-party identity provider services for single sign-on. This allows users to sign on to PingFederate partner applications with popular social platforms such as LinkedIn, Google, or Facebook.

For a current list of integrations, go to the Ping Identity *Integration Directory*.

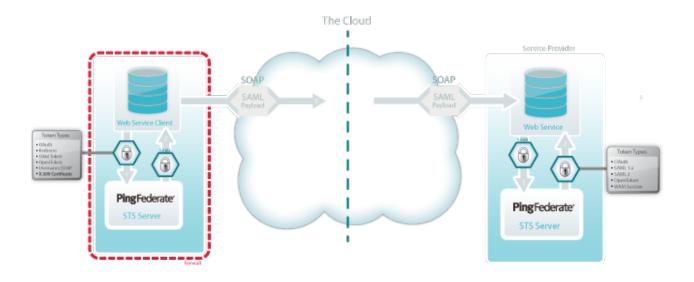
For integrations documentation, go to *Integrations Overview*.

Security token service

The PingFederate WS-Trust Security Token Service (STS) allows organizations to extend single sign-on (SSO) identity management (IdM) to web services.

The STS shares the core functionality of PingFederate, including console administration, identity and attribute mapping, and certificate security management. With PingFederate, web services identify the end user who has initiated a transaction across domains, providing enhanced service while simultaneously ensuring appropriate information access and regulatory accountability. For information about WS-Trust and the role of an STS, see *Web services standards* on page 47.

You can use PingFederate in many different scenarios to address different identity and security problems as they relate to web services, service-oriented architecture (SOA), and Enterprise Service Buses. All of these scenarios share a recommended architectural approach that uses a SAML assertion as the standard security token shared between security domains. For more information, see *About WS-Trust STS* on page 71.



WS-Trust Security Token Service SSO

OAuth authorization server

PingFederate can act as an OAuth authorization server (AS), allowing a resource owner to grant authorization to a client requesting access to resources protected by a resource server (RS).

The OAuth AS issues tokens to clients on behalf of a resource for use in authenticating a subsequent API call—typically, but not exclusively a REST API. The PingFederate OAuth AS issues tokens to clients in several different scenarios, including:

- A web application wants access to a protected resource associated with a user and needs the user's consent.
- A native application client on a mobile device or tablet wants to connect to a user's online account and needs the user's consent.
- An enterprise application client wants to access a protected resource hosted by a business partner, customer, or software as a service (SaaS) provider.

For information about OAuth and the role of an AS, see OAuth 2.0 and PingFederate AS.

You can configure the PingFederate OAuth AS independently or in conjunction with security token service (STS) and browser-based single sign-on (SSO) for either an identity provider (IdP) or a service provider (SP) deployment. For more information, see *About OAuth* on page 75.



Note:

OAuth AS capabilities might require additional licenses. For more information, contact sales@pingidentity.com.

User account management

Typically, the identity provider (IdP) repository maintains user accounts in an identity federation. However, a service provider (SP) often has its own set of user accounts, which might not always correspond to IdP users.

The SP might need to establish and maintain parallel accounts for remote single sign-on (SSO) users to enforce authorization policy, customize user experience, comply with regulations, or a combination of such purposes.

PingFederate provides two kinds of user provisioning for browser-based SSO to facilitate cross-domain account management, one designed for an IdP, and one for an SP:

- At an IdP site, an administrator automatically provisions and maintains user accounts for partner SPs who have implement the System for Cross-domain Identity Management (SCIM) or, when using optional plugin software as a service (SaaS) connectors, for selected hosted-software providers...
- At an SP site, an administrator provisions accounts within the organization automatically from SCIMenabled IdPs or usesinformation from SAML assertions received during SSO events.

For more information, see *User provisioning* on page 97.

Enterprise deployment architecture

PingFederate's enterprise-deployment architecture manages all protocol definitions, public key infrastructure (PKI) keys, policies, profiles, and so forth in a single location, eliminating the need to maintain redundant copies of these configurations and trust relationships.

When you need to add new protocols, profiles, or use cases, you configure them once to make them available to your entire organization.

PingFederate improves security by creating a single "doorway" through which all identity information must travel regardless of who the users are or in which direction they travel. Internal users accessing external applications and external users accessing internal systems both use the same doorway.

The single-doorway approach provides 100% visibility for all federation activities. PingFederate's extensive auditing and logging capabilities enable you to satisfy all of your logging-related compliance and servicelevel requirements from a single location, instead of needing to acquire and consolidate disparate logs from throughout your organization.

PingFederate reduces complexity and learning curves by providing a single configuration model supporting different protocols. The administrative console minimizes the potential for errors by guiding administrators through configuration steps applicable only to the business use cases they need to support.



For identity providers (IdPs), PingFederate provides connection templates to automatically configure many steps in the administrative console for several use cases, including setting up SSO connections to selected SaaS vendors. For more information, see Outbound provisioning for IdPs on page 97.

Additional features

PingFederate's lightweight, standalone architecture allows its server to integrate and coexist with existing home-grown and commercial identity management (IdM) systems and applications to provide the benefits of standards-based single sign-on (SSO) and API security integration without the cost and complexity of deploying a complete IdM system.

PingFederate provides a suite of quick-install integration kits configured from within the PingFederate administrative console to complete the first- and last-mile integration with your existing IdM systems and web applications. Download PingFederate integration kits from the Ping Identity Downloads website.



Multiple security-domain, multi-protocol federation

Integration kits enable rapid session integration with both existing authentication services and target applications. PingFederate also includes a Software Development Kit (SDK) for creating custom integrations. For more information, see SSO integration kits and adapters on page 83.

Token translators

Ping Identity offers special token processors for an IdP and token generators for an SP to enable the WS-Trust security token service (STS) to validate and issue a variety of token types. These plug-ins supplement built-in SAML token processing and generation and handle the local identity tokens required in many security contexts. For more information, see *Token processors and generators* on page 72.

SaaS connectors

SaaS connectors offer a streamlined approach for browser-based SSO to selected SaaS providers, including automatic user provisioning and deprovisioning. The connector packages include quick-connection templates, which automatically configure endpoints and other connection information for each provider. For more information, see *Outbound provisioning for IdPs* on page 97.

Cloud identity connectors

Ping Identity offers social identity integration with social networking sites, including Google, Yahoo!, Twitter, LinkedIn, and Facebook. Connectors leverage OpenID 2.0 or user logins for registration and access to cloud-based applications.

About PingOne

PingOne for Enterprise is a cloud-based identity as a service (IDaaS) framework for secure identity access management. Integrating PingOne for Enterprise with PingFederate provides a powerful solution combining the benefits of an on-premise deployment with the flexibility of a cloud solution.

For more information on PingOne, see *PingOne for Enterprise*.

Key concepts

This section provides background information and preparation to help administrators understand and use PingFederate.

Connection Types

PingFederate features an integrated administrative console for configuring connections to identity-federation partners. The four connection types include:

 Browser-based single sign-on (SSO) – Called Browser SSO in the administrative console, this term refers to standards-based secure SSO, which generally depends on a user's browser to transport

- WS-Trust security token service (STS) Employs the PingFederate STS, which enables web service
 clients (WSCs) and web service providers (WSPs) to extend SSO to identity-enabled web services at
 provider sites. For more information, see the *About WS-Trust STS* on page 71. These standards,
 including WS-Trust, do not rely on the user's browser for message transport.
- OAuth Assertion Grant Exchanges a SAML assertion or a JSON Web Token for an OAuth access token with the PingFederate authorization server (AS). For more information, see *About OAuth* on page 75.
- Provisioning Provides automated cross-domain inbound and outbound user management. For more information, see *User provisioning* on page 97.

You can configure the types of connections together for the same partner or independently.

WS-Trust STS

PingFederate WS-Trust STS allows organizations to extend SSO identity management (IdM) to web services. For more information see, *About WS-Trust STS*.

OAuth

You can configure PingFederate to act as an OAuth authorization server (AS), allowing a resource owner to grant authorization to an OAuth client requesting access to resources hosted by a resource server (RS). For more information, see *About OAuth*.

SSO integration kits and adapters

PingFederate provides bundled and separate integration kits that include adapters that plug into the PingFederate server and agent toolkits that interface with local IdM systems or applications as needed. For more information, see SSO integration kits and adapters.

Security infrastructure

PingFederate security infrastructure supports encrypted messaging, certificates, and digital signing. For more information, see *Security infrastructure*.

Hierarchical plugin configuration

PingFederate allows you to use a configuration of an adapter, as well as certain other PingFederate plugins, as a parent instance from which you can create child instances. For more information, see *Hierarchical plugin configurations*.

Identity mapping

PingFederate enables identity mapping between domains for browser-based SSO and WS-Trust STS. For more information, see *Identity mapping*.

User attributes

Federation transactions require the transmission of a unique piece of information that identifies the user for identity mapping between security domains. For more information, see *User attributes*.

User provisioning

PingFederate provides cross-domain user provisioning and account management. For more information, see *User provisioning*.

PingFederate empowers administrators to deliver a secure and easy-to-use customer authentication, registration, and profile management solution. For more information, see *Customer identity and access management*.

Federation hub use cases

As a federation hub, PingFederate can bridge browser-based SSO between IdPs and SPs. For more information, see *Federation hub use cases*.

Federation planning

An essential first step in establishing an identity federation involves discussions and agreements between you and your connection partners. For more information, see the *Federation planning checklist*.

About WS-Trust STS

The PingFederate WS-Trust Security Token Service (STS) allows organizations to extend single-sign on (SSO) identity management (IdM) to web services.

You can configure the WS-Trust STS for partner connections independently or in conjunction with browser-based SSO for either an identity provider (IdP) or a service provider (SP) deployment. The STS is bundled with separate plug-ins for standard SAML token processing and generation.

For information about WS-Trust and the role of an STS, see Web services standards on page 47.

Connection-based policy

PingFederate employs a partner-connection configuration for both IdP and SP roles, which enables the association of web services authentication policies with federation partners. For more information, see *Connection-based policy*.

Token processor and generator

PingFederate provides support for a variety of security-token formats through token processors and generators. For more information, see *Token processors and generators*.

WSC and WSP support

Ping Identity provides the Java client software development kit (SDK) for enabling web service applications to interact with the PingFederate STS. For more information, see *WSC* and *WSP* support.

STS OAuth integration

PingFederate STS provides several ways to facilitate the use of issued tokens with an OAuth authorization server (AS). For more information, see *STS OAuth integration*.

Connection-based policy

For both the identity provider (IdP) and service provider (SP) roles, PingFederate employs a partnerconnection configuration, which enables the association of web services authentication policies with federation partners.

For Security Token Service (STS) processing, these policies define configurations for handling WS-Trust requests and transferring identity information between security domains. For more information, see *Web services standards* on page 47.

IdP configuration

Use the administrative console in an IdP role to configure WS-Trust request-processing policy for your SP partner including:

- The type of SAML token to create in response to an issue request from a web service client (WSC) application
- The mapping of attributes to include within the issued SAML token
- The key used to create a digital signature for the issued SAML token

SP configuration

Use the administrative console in an SP role to configure WS-Trust request-processing policy for your IdP partner including:

- Whether to validate the incoming SAML token only, or to validate the incoming token and also issue a local token
- The mapping of attributes to include in the locally issued token when applicable
- The certificate used to verify the digital signature for the incoming SAML token
- The key used to decrypt the incoming SAML token when needed

Token processors and generators

PingFederate provides support for a variety of security-token formats through token processors and generators.

These token processors and generators plug into the PingFederate server and deploy similarly to browserbased single sign-on (SSO) adapters. For more information, see SSO integration kits and adapters on page 83.

For an identity provider (IdP), token processors provide a mechanism through which PingFederate can validate an incoming token from a web service client (WSC) and map attributes to be included in the issued SAML token.

For a service provider (SP), token generators provide a mechanism through which PingFederate can generate a local token based upon the incoming SAML token from a WSP and map attributes to be included in that token.

PingFederate only generates SAML 1.1 or 2.0 token when it is configured as an IdP for sending across trust boundaries to a federate SP partner. PingFederate only accepts SAML tokens when configured as an SP. Token plug-ins allow a modular approach for validating and producing the token types used by different applications or systems within a conceptual trust domain. PingFederate provides bundled and separately available token plug-ins.



For direct security token service (STS) token exchange within the same domain or trust boundary, use the PingFederate STS to exchange one token type for another directly, without generating a transitional SAML token. For more information, see *Token translator mappings* on page 377.

PingFederate allows you to use a configuration of a token processor or generator as a parent instance from which you create child instances. For more information, see Hierarchical plugin configurations on page 90.

Bundled token plug-ins

PingFederate comes installed with token processors for an IdP configuration that accept and validate SAML 1.1 or 2.0 tokens, OAuth bearer access tokens, JSON web tokens (JWT), username tokens, and Kerberos tokens. For more information, see *Token models and management* on page 75). SAML tokens are issued on the IdP side through built-in browser-based SSO capabilities.

For an SP configuration, PingFederate provides token generators for issuing local SAML 1.1 or 2.0 tokens. PingFedearate validates incoming SAML tokens using built-in capabilities.

Ping Identity provides token plug-ins called token translators to work with various authentication systems and identity management (IdM) systems. You can download the available plug-ins from the Downloads website.

WSC and WSP support

Ping Identity provides the Java client software development kit (SDK) for enabling web service applications to interact with the PingFederate security token service (STS).

For web service client (WSC) STS clients, PingFederate provides built-in protocol support for Windows Identity Foundation (WIF) applications based on the Windows Communication Foundation (WCF) framework.



Note:

The WIF framework includes WS-* protocol support and can interact natively with PingFederate.

Client SDK

The STS Java client SDK provides interfaces that create the WS-Trust Request Security Token (RST) and Request Security Token Response (RSTR) messaging to interact with the PingFederate STS endpoints. Using the SDK library, applications are not responsible for forming these WS-Trust protocol messages, and instead interact only with the tokens themselves.

The SDK is available for download on the Ping Identity Downloads website.

Windows Identity Foundation clients

PingFederate natively supports STS clients using claims-based WIF technology. Claims-based federated identity for web services is a part of the WS-Trust standard that permits client applications to make accesspolicy decisions, when specifically categorized user attributes are sent in the security token. For more information, see *Attribute contracts* on page 93.

The PingFederate STS supports the following bindings in the .NET federated-security scenarios with WS-Trust:

- WSFederationHttpBinding
- WS2007FederationHttpBinding

Additionally, the PingFederate STS supports the following bindings for RST and RSTR interactions with .NET. Support for these bindings is limited to the Username, x509, SAML 1.1, and SAML 2.0 token types:

- WSHttpBinding
- WS2007HttpBinding



Note:

For token types such as Kerberos, where customizing default bindings might be necessary, the PingFederate STS supports the use of customBinding.

For more information about bindings, see Microsoft's System-Provided Bindings.

To expedite configuring their applications, PingFederate provides metadata for developers. PingFederate offers two varieties of metadata, which work together to arrive at functional WSC and web service provider (WSP) configurations:

 STS Metadata Exchange at /pf/sts mex.ping, which contains connection details relating to the SP partner.

For more information about claim-based federated identity, see Microsoft's *A Guide to Claims—based Identity and Access Control*.

STS OAuth integration

PingFederate security token service (STS) provides several ways to facilitate the use of issued tokens with an OAuth authorization server (AS).

OAuth token processor

This token processor provides a mechanism through which PingFederate STS can validate an incoming OAuth Bearer access token. The token processor reads and validates the access token and returns any additional user attributes defined.

JWT bearer token grant type

```
urn:ietf:params:oauth:grant-type:jwt-bearer
```

This token request returns a JSON Web Token (JWT) that a web service client (WSC) can use to request OAuth access tokens from any OAuth AS that supports using JWTs as authorization grants, as defined in JSON Web Token (JWT) Profile for OAuth 2.0 Client Authentication and Authorization Grants specification.

OAuth access token with JWT bearer token grant type

```
\verb|oauth-v2:access:token:response| \verb|via|| urn:ietf:params:oauth:grant-type:jwt-bearer| \\
```

This proprietary token request is similar to the JWT Bearer Token grant type but returns an OAuth access token directly. Acting as an identity provider (IdP), PingFederate generates the intermediate JWT and requests an access token from the OAuth AS on behalf of the WSC. The AS endpoint is obtained from the AppliesTo element of the WS-Trust request security token (RST) message.

SAML 2.0 bearer assertion grant type

```
urn:ietf:params:oauth:grant-type:saml2-bearer
```

This token request returns an encoded SAML assertion that a WSC can use to request OAuth access tokens from any OAuth AS that supports the SAML 2.0 Profile for OAuth 2.0 Client Authentication and Authorization Grants specification.

OAuth access token with SAML 2.0 bearer assertion grant type

```
oauth-v2:access:token:response|via|urn:ietf:params:oauth:grant-type:saml2-bearer
```

This proprietary token request is similar to the SAML 2.0 Bearer Assertion grant type but returns an OAuth access token directly. Acting as an IdP, PingFederate generates the intermediate, encoded SAML assertion and requests an access token from the OAuth AS on behalf of the WSC. The AS endpoint is obtained from the AppliesTo element of the WS-Trust RST message.

These capabilities bridge the WS-Trust client-STS relationship and the trust relationship the same client with an OAuth AS, allowing the client to obtain additional resources on behalf of already-authenticated users in follow-on transactions.

Configuring PingFederate as an OAuth authorization server (AS) allows a resource owner (RO), typically an end user, to grant authorization to an OAuth client requesting access to the resource server (RS).

The OAuth AS issues tokens to clients on behalf of an RO for use in authenticating a subsequent API call to the RS, typically, but not exclusively, a REST API call.



If your PingFederate license does not include the OAuth AS capabilities, contact sales@pingidentity.com.

You can configure the PingFederate OAuth AS independently or in conjunction with security token service (STS) or browser-based single sign-on (SSO) for either an identity provider (IdP) or a service provider (SP) deployment.

In an IdP deployment, an IdP adapter is used to authenticate and provide user information for the access token. In an SP deployment, the inbound SAML assertion is used to provide authentication information about the user associated with the access token through an OAuth attribute mapping in the IdP connection.

For an STS IdP, PingFederate provides an OAuth token processor that validates incoming OAuth Bearer access tokens.

Delegated access types

To enhance access control, PingFederate supports both explicit and implicit delegation of transaction approval.

Explicit delegation

This is the most common OAuth use case, which involves a resource owner (RO) who explicitly delegates the authority to a client to make API calls to a resource server (RS) and is asked to approve the transaction. This is the type of delegation inherent in web redirect flow.

Implicit delegation

Implicit delegation also generally involves a client who calls an API on behalf of a user. However, the client's authority is implied by the nature of the transaction, and the user is not specifically asked to approve the transaction.

Token models and management

Successful OAuth transactions require an OAuth authorization server (AS) to issue tokens characterized by both security model and data model for use in authenticating an API call.

Token security model

A token security model refers to the conditions that must be met by a client in order to use a token on an API call. The currently supported model is a Bearer Token. A client's presentation of the token – for example, as a parameter on the API call - to the resource server (RS) is interpreted as providing sufficient proof to the RS that the client received the same token from the OAuth AS.

Token data model

A token data model refers to whether the token carries identity and security information or acts as a pointer to the information.

Self-contained tokens (JSON Web Tokens)

Contain identity and security information and attributes in a transport format such as JSON, signed by the AS and verified directly by the RS.

Reference tokens (Internally Managed Reference Tokens)

Token management

PingFederate supports multiple access token management instances, providing flexibility for enterprises where deployments require different token data models, token lifetimes, attribute contracts, token validation rules, or any combination of them for various clients.

Grant types

To obtain an access token, a client interacts with an OAuth authorization server (AS), sending a request for an access token that includes an access grant. An access grant is also used when a resource server (RS) requests validation of an access token from the AS.

Primary grant types

OAuth defines several different access grant types. Each grant type reflects different authorization mechanisms.

Authorization code

```
authorization code
```

An authorization code is returned to the client through a browser redirect after the resource owner (RO) gives consent to the AS. The client subsequently exchanges the authorization code for an access token and often a refresh token. RO credentials are never exposed to the client.

Resource owner password credentials

```
password
```

The client collects the RO's password and exchanges it at the AS for an access token and often a refresh token. This grant type is suitable in cases where the RO has a trust relationship with the client, such as its computer operation system or a highly-privileged application because the client must discard the password after using it to obtain the access token.

Refresh token

```
refresh_token
```

A refresh token often returns with an access token. Once the original access token expires, the corresponding refresh token sends to the AS to obtain a fresh access token without requiring the RO to re-authenticate. This allows short-lived access tokens to exist between the client and the resource server and long-lived tokens between the client and the AS.

The refresh token grant type only works in conjunction with either the authorization code or RO password credentials grant type.

Implicit

```
implicit
```

A browser redirect responds to the RO authorization request, rather than an intermediate authorization code, and returns an access token to the client. This grant type works for clients incapable of keeping client credentials confidential for use in authenticating with the AS, such as client applications implemented in a browser using a scripting language like JavaScript.

Client credentials

```
client_credentials
```

The client presents its own credentials to the AS in order to obtain an access token. This access token is either associated with the client's own resources, and not a particular RO, or with a RO for whom the client is otherwise authorized to act.

Extension grant types

OAuth provides an extension mechanism for defining new extension grant types to support additional clients or to provide a bridge between OAuth and other trust frameworks. An OAuth client uses an extension grant type by specifying an absolute URI as the value of the grant type parameter and by adding any additional parameters necessary when contacting the token endpoint at /as/token.oauth2.

PingFederate supports the following extension grant types:

Assertion grants

JWT Bearer

```
urn:ietf:params:oauth:grant-type:jwt-bearer
```

The client obtains a JSON web token (JWT) and uses it to request an access token from the AS. This grant type allows a client to use an existing trust relationship, expressed through a JWT, without a direct user approval step at the AS.

SAML 2.0 Bearer

```
urn:ietf:params:oauth:grant-type:saml2-bearer
```

The client obtains a SAML 2.0 Bearer assertion and uses it to request an access token from the AS. Similar to the JWT Bearer grant type, this grant type allows a client to use an existing trust relationship, expressed through a SAML assertion, without a direct user approval step at the AS.



Note:

The SAML assertion used for this grant type generally cannot be a browser-based SSO assertion. To ensure its validity, the assertion must be associated with WS-Trust security token service (STS) processing.

Client-initiated backchannel authentication (CIBA) grant

```
urn:openid:params:grant-type:ciba
```

The client presents an identity hint to the AS. The AS identifies the RO based on the hint provided and then authenticates and obtains authorization from the RO through an out-of-band flow. Depending on the setup, the client either polls the AS for the authorization result or wait for a signal from the AS to return to the AS for the authorization result. If the RO approves the authentication request, the AS returns an access token to the client. Otherwise, the AS returns an error message per the specification.

Device authorization grant

```
urn:ietf:params:oauth:grant-type:device code
```

The client presents a device code and user code to the AS in order to identify the deviceauthorization session and obtain an access token. This access token is associated with a RO for whom the client is otherwise authorized to act.

token exchange grant

```
urn:ietf:params:oauth:grant-type:token-exchange
```

The client presents a security token to the AS. In exchange, the AS returns another kind of security token to the client. The new token might be an access token that is more narrowly scoped for a downstream service or it could be an entirely different kind of token. This grant type supports subject and actor tokens employing impersonation and delegation.

Validation grant

```
urn:pingidentity.com:oauth2:grant type:validate bearer
```

This proprietary PingFederate OAuth extension enables an RS to act as a client in the request/ response exchange with PingFederate as the AS in this scenario. The grant type allows an RS to check with PingFederate on the validity of a bearer access token received from a client making a protected-resources call.

Scopes

In addition to OAuth, PIngFederate supports the use of scopes to constrain and define access privileges.

OAuth provides a mechanism to constrain the privileges associated with an access token, whereas scopes provide a way to more specifically define the privileges requested and granted. Generally, a client specifies the desired scopes when sending an authorization request to the authorization server. If the user approves, the authorization server issues an access token with these scopes.

Scopes are configured globally using the System # OAuth Settings # Scope Management configuration wizard. Once defined, you can manage the availability of scopes on a client-by-client basis.

Static scopes and dynamic scopes

As an authorization server, PingFederate supports the concepts of static scopes and dynamic scopes. To define a static scope, use a text value such as read bank account. To define a dynamic scope, use a text value with a variable component represented by a wildcard, such as read bank account txn:*. As illustrated, dynamic scopes allow clients to request authorization using scope values with a variable component from one request to another.

Consent approval

With Authorization Code, Implicit, and Device Authorization grant types, an authorization server (AS) prompts the user to grant authorization to share user information. Once granted, the AS issues an access token to the client who uses it to access information from the resource server (RS).

Default consent user interface

PingFederate handles the consent approval process by presenting the Request for Approval window to the user by default. This window displays a list of requested permissions, scopes, along with their descriptions as configured in PingFederate. It is up to the user to approve or deny individual scopes.

External consent user interface

As use cases evolve towards giving users more control over their data, it becomes more important to provide detailed information about the requests. While the scope description can help, PingFederate also supports the use of an external web application to prompt for authorization consent. This approach opens up the opportunity to retrieve additional information specific to the users. For example, the web application can be written in such a way that when a client requests the read bank account scope, the web application retrieves the user's customer information file and gives the user the ability to choose which accounts to make available to the client.

Client management and storage

OAuth clients interact with an authorization server (AS) to obtain access tokens and optionally refresh tokens to access protected resources on resource servers.

PingFederate provides administrators the flexibility to manage OAuth clients using the following interfaces:

- The administrative console
- The administrative API
- The OAuth Client Management Service

Additionally, PingFederate supports dynamic client registration based on the OAuth 2.0 Dynamic Client Registration Protocol specification.

Storing client records in XML files by default allows administrators to manage clients using the administrative console and the administrative API. It also allows developers to submit client creation Alternatively, because the OAuth Client Managment Service requires external storage of client records, PingFederate supports configuration to store client records externally on a database server, a directory server, or some other storage medium through the use of the PingFederate SDK. Under this configuration, the configuration archive does not include client records.

Client authentication schemes

Most OAuth and OpenID Connect use cases require the client application to authenticate successfully before its requests can be processed further.

As an OAuth authorization server (AS), PingFederate supports the following client authentication schemes:

- Client secret for HTTP Basic authentication
- Client TLS certificate for mutual TLS authentication
- Private key JWT for the private_key_jwt client authentication method, as defined in the OpenID Connect specification
- None when authentication is not required

When deployed as an OpenID Connect Relying Party (RP), PingFederate authenticates through client secret and private key JSON web tokens (JWT). It also handles the scenario where authentication is not required.

Dynamic client registration

As an OAuth provider, PingFederate supports a number of OAuth protocols.

PingFederate supports dynamic client registration based on the OAuth 2.0 Dynamic Client Registration Protocol specification. When enabled, it allows developers to register OAuth clients through an API based on open standards.

Transient grants and persistent grants

There are two types of OAuth authorization grants: transient grants and persistent grants.

Transient grants

Transient grants are valid only for the lifetime of their respective access tokens. Transient authorizations include those obtained by OAuth clients in the following manners:

- Grants obtained by using the authorization code, resource owner credentials, or device authorization grant type, without the refresh token grant type
- Grants obtained by using client credential, JWT bearer, SAML 2.0 bearer assertion, or token exchange grant type

Transient grants are not preserved.

Persistent grants

Persistent grants typically bear a longer lifetime than their respective access tokens do. Persistent authorizations include those obtained by OAuth clients in the following ways:

Grants obtained or updated using the authorization code, resource owner credentials, or device authorization grant type, in conjunction with the refresh token grant type



Note:

If the use cases involve mapping attributes from authentication sources, such as IdP adapter instances or IdP connections, or password credential validator (PCV) instances to the access tokens, directly or through persistent grant-extended attributes, storing these attributes from authentication sources

and their values along with the persistent grants maintains them for reuse when clients subsequently present refresh tokens for new access tokens.

Grants obtained or updated by using the implicit grant type, for which PingFederate is configured to reuse existing persistent grants



Note:

If the use cases involve mapping attributes from authentication sources or PCV instances to the access tokens, runtime procedures obtain attribute values for each token request, but persistent grants do not store with attributes or their values.

Persistent grant lifetime and maintenance

Persistent grants, and any associated attributes and their values, remain valid until the grants expire or until PingFederate explicitly revokes them or cleans them up.

Grants persist without any expiration information. Grants also persist with an idle timeout window, a maximum lifetime, or both. If you configure an idle timeout value, the idle timeout window slides when a persistent grant updates. When you have an idle timeout value configured without a maximum lifetime, persistent grants remain valid until they expire due to inactivity or until the grant storage revokes or removes them. When you have an idle timeout value configured with a maximum lifetime, persistent grants remain valid until they expire due to inactivity or lifetime expiration or until the grant storage removes them.

PingFederate removes expired grants and the associated attributes from the grant datastore once a day. The frequency and the size of the cleanup batch are configurable. Optionally, PingFederate caps the number of persistent grants on a basis of the combination of user, client, and grant type.

Persistent grant storage

Support for persistent grants requires the use of a database server or a directory server for long-term storage. PingFederate also supports other storage solutions through the PingFederate SDK. For more information, see OAuth grant datastores on page 889.

PingFederate uses a built-in HSQLDB database as its grant datastore after the initial setup.



CAUTION:

Use the built-in HSQLDB only for trial or training environments. For testing and production environments, always use a secured external storage solution for proper functioning in a clustered environment.

Testing involving HSQLDB is not a valid test. In both testing and production, it might cause various problems due to its limitations and HSQLDB involved cases are not supported by Pingldentity.

Grant storage and management

PingFederate uses a built-in HSQLDB database as its persistent grant datastore after the initial setup.



CAUTION:

Use the built-in HSQLDB only for trial or training environments. For testing and production environments, always use a secured external storage solution for proper functioning in a clustered environment.

Testing involving HSQLDB is not a valid test. In both testing and production, it might cause various problems due to its limitations and HSQLDB involved cases are not supported by Pingldentity.

Persistent grants, and any associated attributes and their values, remain valid until the grants expire or until PingFederate explicitly revokes them or cleans them up.

For revocation, PingFederate provides two endpoints.

Token revocation endpoint

The token revocation endpoint allows clients to notify the authorization server that they no longer need a previously-obtained refresh or access token. The revocation request invalidates the actual token and possibly other tokens based on the same authorization grant.

Grant-management endpoint

The grant-management endpoint allows resource owners to view and optionally revoke the persistent access grants they have authorized.

Intended for OAuth clients, the token revocation endpoint is the endpoint to which clients send their token revocation requests. The grant-management endpoint is for resource owners. It displays a list of grants the resource owners have made. Resource owners can view and optionally revoke one or more grants as they see fit.

Mapping OAuth attributes

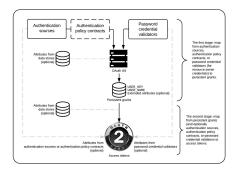
Mapping OAuth attributes is a two-stage processing workflow.

The two stages of mapping OAuth attributes include:

- The first stage: map attributes from authentication sources (such as IdP adapter instances and IdP connections), authentication policy contracts, and Password Credential Validator instances (for resource owner credentials) to persistent grants.
- The second stage: map from persistent grants, authentication sources, authentication policy contracts, authentication context, or Password Credential Validator instances to access tokens.



This two-stage mapping workflow is different from other mapping scenarios in PingFederate, which involve just a one-phase configuration.



The first stage

To accomplish the first stage, setting up persistent grants requires mapping, including a user key and all extended attributes.

The mappings use attributes obtained during initial authentication events within PingFederate, namely attributes from IdP adapter instances, attributes from assertions through IdP connections, attributes from authentication policy contracts, or attributes returned by password credential validator instances. Configure data store queries to accomplish different tasks, such as retrieving the user identifier from an LDAP directory server as the user key.

Important:

The USER KEY attribute values must be unique across all end users, because the USER KEY attribute is the user identifier to store and to retrieve persistent grants. For example, the samaccountName attribute value of an end user in one domain might match that of another end user in another domain. In this case, you can map the Subject DN attribute to the USER KEY attribute.

The second stage

The second mapping configuration involves mapping from persistent grants the user keys, any extended attributes derived from the first stage, or both into OAuth access tokens.

When the authentication context matches specific mappings between the authentication sources, authentication policy contracts, or password credential validators, the attributes map into access tokens. The HTTP RequestJava object retrieves the authentication method that a client uses, or the private key JWT for client authentication if the client uses the private_key_jwt authentication method. The HTTP Request Java object maps it into the access tokens.

Data stores used here retrieve any required user attributes.

Runtime processing

At runtime, the first time a client requests an OAuth token, the process employs the two mapping sequences. Every time an existing persistent grant requests a new access token, the process invokes the second mapping.

OAuth user-facing windows

The PingFederate OAuth authorization server (AS) presents five windows to end users during OAuth transactions.

The windows include:

- One prompting for approval of the requested scopes
- One providing a means of revoking persistent access grants
- Three others displaying information for the purpose of connecting OAuth-capable internet of things (IOT) based on the OAuth 2.0 Device Authorization Grant specification

Administrators can customize these windows as needed.

OpenID Connect

As an extension of OAuth capabilities, PingFederate supports an optional configuration for OpenID Connect, a modern protocol for secure, lightweight transfer of authentication and user attributes.

OpenID Provider support

As an OpenID Provider (OP), PingFederate supports both the Basic Client and Implicit Client profiles defined in the standard. In both profiles, the end result releases an ID token and an OAuth access token; however, depending on associated grant types, PingFederate might also release a refresh token.

The ID token is an integrity-secured, self-contained token in JSON Web Token (JWT) format containing claims about the user. A client uses the ID token to identify the user accessing the client application through an OP. A client may subsequently use the OAuth access token to retrieve additional claims about the user, such as a complete profile containing full name, email, phone, and other schema elements defined in an OpenID Connect policy from the Userinfo endpoint (/idp/userinfo.openid).

For session management, PingFederate provides a front-channel endpoint for OAuth clients using the OpenID Connect protocol to close other associated sessions at /idp/startSlO.ping and a back-channel web service for clients to revoke end-user sessions at /pfws/rest/sessionMamt/ revokedSris.

Relying Party support

As a Relying Party (RP), PingFederate is capable of leveraging identities from OPs to complete browser-based SSO requests. In this use case, PingFederate is the requesting OAuth client application.

The setup involves establishing an IdP connection to the OP. PingFederate retrieves identity information from the OP and passes the end-user claims, which are user attributes in an ID token, to one or more target applications. This configuration allows administrators to take advantage of their existing last-mile integration and expand the horizon of their applications to additional partners using the OpenID Connect protocol.

PingFederate is also capable of sending request parameters through self-contained, signed JWTs, thus adding a layer of security to the transmission of the request parameters. Additionally, if the ID token contains a state hash, PingFederate validates it.

CORS support for OAuth endpoints

PingFederate supports cross-origin resource sharing (CORS) for several OAuth endpoints.

The supported OAuth endpoints include:

- /as/token.oauth2
- /as/revoke token.oauth2
- /idp/userinfo.openid
- /pf-ws/rest/oauth/grants/
- /pf/JWKS
- /.well-known/openid-configuration
- /as/bc-auth.ciba

As needed, administrators can add or remove allowed origins using the administrative console on the **Authentication Application** page. For more information, see *Configuring authentication applications* on page 393. Once configured, client-side web applications from the trusted origins are allowed to make requests to the PingFederate authorization server for the purpose of accessing protected resources, such as obtaining or renewing access tokens with refresh tokens, presenting access tokens for revocation, querying additional claims (user attributes), and retrieving OpenID Provider configuration information and JSON Web Key Sets.

SSO integration kits and adapters

As a standalone server, PingFederate must be programmatically integrated with end-user applications and identity management (IdM) systems to complete the "first- and last-mile" implementation of a federated identity network for browser-based single-sign-on (SSO). Documentation for integration kits is available on the Ping Identity *website*.

For an identity provider (IdP), the first mile of this integration process involves providing a mechanism through which PingFederate looks up a user's current authenticated session data, such as a cookie, or authenticates a user without such a session. For a service provider (SP), the last mile involves enabling PingFederate to supply information needed by the target application to set a valid session cookie or other application-specific security context for the user. To enable both sides of this integration, PingFederate provides bundled and separately available integration kits, which include adapters that plug into the PingFederate server and agent toolkits that interface with local IdM systems or applications, as needed. In addition, PingFederate provides plugin authentication selectors, which enable dynamic selection of authentication sources based on administrator-specified criteria. For more information, see the PingFederate SSO integration overview on page 60.

Bundled adapters

PingFederate comes bundled with a set of adapters.

Identifier First Adapter

When a variety of user types authenticate at PingFederate, it is often better to ask the user for their identifier first, determine their user population, and prompt the user with the desired authentication requirements and experience. The Identifier First Adapter is designed to handle this use case. For more information, see *Identifier First Adapter* on page 305.

HTML Form Adapter

Used in conjunction with Password Credential Validators. These adapters provide integration with user-datastores in directory servers or locally. For more information, see *HTML Form Adapter* on page 287.

Kerberos Adapter

Provides a seamless desktop SSO experience for Windows environments and supports authentication mechanism assurance from the Active Directory domain service. For new configurations and as a simpler alternative to the separately-available IWA Integration Kit, use this adapter. For more information, see *Kerberos Adapter* on page 313.

OpenToken Adapter

Provides a generic interface for integrating with various applications, including Java- and .NET-based applications. For more information, see *OpenToken Adapter* on page 318.

Composite Adapter

Allows multiple configured IdP adapters to execute in sequence. Depending on the authentication context, use this capability, called adapter chaining, for either single-adapter usage or to support multifactor authentication through a series of adapters. For more information, see *Composite Adapter* on page 283.

HTTP Basic Adapter

Used in conjunction with Password Credential Validators. These adapters provide integration with user-data stores in directory servers or locally. For more information, see *HTTP Basic Adapter* on page 303.

PingID

PingID is a cloud-based authentication service that binds user identities to their devices, making it an effective multifactor authentication solution. For more information, see *PingID documentation*.

Bundled authentication selectors

PingFederate provides plugin authentication selectors, which enable dynamic selection of authentication sources based on administrator-specified criteria. Along with the Composite Adapter and token authorization, the selectors enable dynamic integration with an organization's authentication or authorization policies, also known as adaptive federation.



Tip:

To select subsequent selectors or authentication sources for handlingcomplex hierarchical accesspolicy decisions, use the results of authentication-selection criteria evaluation. For more information, see *Authentication policies* on page 220.

CIDR Authentication Selector

Provides a means of choosing authentication sources or other authentication sources at runtime based on whether an end-user's IP address falls within specified ranges using Classless Inter-Domain Routing notation. This selector allows administrators to determine, for example, whether an SSO request originates inside or outside the corporate firewall and use different authentication integration accordingly. For more information, see *Configuring the CIDR Authentication Selector* on page 222.

Cluster Node Authentication Selector

Provides a means of picking authentication sources or other authentication sources at runtime based on the PingFederate cluster node that is servicing the request. For example, you can configure this selector to choose whether PingFederate attempts Integrated Windows Authentication based on the PingFederate cluster node with which a Key Distribution Center is associated. For more information, see *Configuring the Cluster Node Authentication Selector* on page 223.

Connection Set Authentication Selector

Provides a means of selecting authentication sources or other authentication sources at runtime based on a match found between the target SP connection used in an SSO request and SP connections configured within PingFederate. For example, administrators with different requirements for SP connections can override connection adapter selection on an individual connection basis. For more information, see *Configuring the Connection Set Authentication Selector* on page 224.

Extended Property Authentication Selector

Enables PingFederate to choose configured authentication sources or other selectors based on a match found between a selector result value and an extended property value from the invoking browser-based SSO connections or OAuth client. For more information, see *Configuring the Extended Property Authentication Selector* on page 225.

HTTP Header Authentication Selector

Provides a means of choosing authentication sources or other authentication sources at runtime based on a match found using wildcard expressions in an HTTP header. This selector allows administrators to determine, for example, authentication behavior based on the type of browser. For more information, see *Configuring the HTTP Header Authentication Selector* on page 227.

HTTP Request Parameter Authentication Selector

Provides a means of selecting authentication sources or other authentication sources at runtime based on query parameter values in the HTTP request. For more information, see *Configuring the HTTP Request Parameter Authentication Selector* on page 229.

OAuth Client Set Authentication Selector

Enables PingFederate to choose configured authentication sources or other selectors based on a match found between the client information in an OAuth request and the OAuth clients configured in the PingFederate OAuth authorization server (AS). This selector allows you to override client authentication selection on an individual client basis in one or more authentication policies. For more information, see *Configuring the OAuth Client Set Authentication Selector* on page 231.

OAuth Scope Authentication Selector

Provides a means of selecting authentication sources or other authentication sources at runtime based on a match found between the scopes of an OAuth authorization request and scopes configured in the PingFederate OAuth authorization server (AS). For example, if a client requires write access to a resource, administrators can configure the selector to choose an adapter that offers a stronger form of authentication such as the X.509 client certificate rather than username and password. For more information, see *Configuring the OAuth Scope Authentication Selector* on page 232.

Provides a means of picking authentication sources or other authentication sources at runtime based on the authentication context requested by an SP, for SP-initiated SSO. Configured authentication sources are mapped either to SAML-specified contexts or any ad-hoc context agreed upon between the IdP and SP partners. For more information, see *Configuring the Requested* AuthN Context Authentication Selector on page 233.

Session Authentication Selector

Enables PingFederate to choose a policy path at runtime based on whether the user already has a PingFederate authentication session for a particular source. For more information, see Configuring the Session Authentication Selector on page 234.



Note:

Authentication selectors rely on HTTP requests, HTTP headers, POST data, or a combination of them. Ensure that standard security measures are in place when using these selectors.

Integration kits

Ping Identity regularly develops and maintains integration kits and adapters to work with applications and leading identity management systems. Download available kits from the Ping Identity Downloads website. Ping Identity adds additional authentication selectors to the download site. Contact sales@pingidentity.com with your specific authentication-selection capabilities.

Software development kit (SDK)

The PingFederate SDK provides a flexible means of creating custom adapters to integrate federated identity management into your system environment. For more information, see the PingFederate SDK Developer's Guide on page 1021.

Security infrastructure

This section describes the PingFederate security infrastructure that supports encrypted messaging, certificates, and digital signing.

PingFederate's configuration windows integrate these functions to provide complete control over certificate generation and authentication verification.

Digital signatures

A digital signature is a way to verify the identity of a person or entity who originates an electronic document and ensure that the message has not been altered.

Both SAML, including security token service (STS) tokens, and WS-Federation electronic documents use digital signatures.

Handling a digital signature involves message signing, signature and certificate validation, and signingpolicy coordination between connection partners.

Message signing

Certificates contain information about the certificate owner along with a public key. Applying a digital signature creates and encrypts a hash from the signing message using the private key.

PingFederate provides a choice of signature encryption algorithms when you require a stronger algorithm.

To ensure the integrity of SAML messages or security token service (STS) tokens, we recommend digital signing practices using public/private keypairs in conjunction with X.509 certificates.



Note:

Ping Identity recommends a certificate signed by a certificate authority (CA); however, PingFederate will work with self-signed or untrusted third-party-signed certificates. After generating a keypair and a selfsigned certificate, use PingFederate to create a certificate signing request (CSR) and send it to a CA for signing. After the CA has generated a CSR, import it into PingFederate's certificate management system. PingFederate's trusted store or the Java runtime cacerts store must contain the CA's certificate.

PingFederate enables signing and validation of requests and responses. Additionally, PingFederate provides for certificate generation, import and export functionality, CSR generation, and application of digital signatures. You have the option to create reusable global signing certificates across your federated connection base and import signature verification certificates for each partner. For more information, see Manage digital signing certificates and decryption keys on page 608.



Ping Identity recommends generating unique certificates for each connection, which limits exposure if the private key becomes compromised.

Signature validation

After receiving a signed message, PingFederate verifies the signature using the public key that corresponds with the private key used to sign the message or token. Verification involves creating a hash of the received message, using the signing partner's public key to decrypt the hash sent with the original message, and verifying that both hash values are equal.

Certificate validation

PingFederate always checks certificates to see if they have expired when they are initially imported. It also checks certificates at runtime when they are used to verify incoming signed assertions.

PingFederate also checks to see whether a certificate has been revoked, using either certificate revocation lists (CRLs) or the online certificate status protocol (OSCP). Depending on the content of the certificate in question and your requirements, the server will perform either of these checks during single sign-on (SSO) or single log-out (SLO) processing for the following cases:

- Signature verification
- Validation of a client certificate used for authentication to PingFederate when the server is handling direct client requests
- Validation of the server SSL certificate when PingFederate acts as the client making an HTTPS request to a separate server

If the system encounters an expired or revoked certificate, the associated SSO or SLO transaction fails at runtime and writes an error to the transaction log. In the administrative console, the Status column of the respective Certificate Management list identifies the expired or revoked certificate.

CRL revocation checking

This process involves guerying a CRL distribution-point URL and ensuring that a certificate is not on the returned revocation list maintained at the site. The certificate specifies the URL.

The administrative console does not need any setup to enable CRL checking. PingFederate automatically checks CRLs under the following conditions:

- The certificate contains the URL where the CA maintains its CRL.
- The URL is accessible.
- The CRL returns signed with a verified signature.
- The OCSP setup does not explicitly disable CRL validations as a failover option.

OCSP exists as an alternative to CRL validation and provides a more centralized and potentially more reliable means of checking certificate status. In this scenario, the incoming certificate embeds an OCSP Responder URL or a configured default URL to query the certificate status.

The primary difference between OCSP and CRL checking is how the verification occurs. CRL checking requires the requesting client to determine if the certificate has been revoked, or if any of the certificates in the chain of issuer certificates has been revoked, based on the returned CRL. With OCSP, the client sends the certificate itself, and the Responder server handles revocation checking to return the certificate status.

As a PingFederate administrator, enable and configure OCSP processing in the administrative console exclusively or in conjunction with CRL checking as a backup.

For more information about OCSP, see tools.ietf.org/html/rfc2560.

For configuration steps, see *Configuring certificate revocation* on page 623.

Digital signing policy coordination

To coordinate digital signature policy, partners must first agree about whether they will sign SAML messages or tokens.

In some cases, such as SAML security token service (STS) tokens and single sign-on (SSO) assertions sent across the POST binding, the protocol specifications require signatures. The PingFederate administrative console and the runtime protocol engine enforce these requirements. Some partner connections specify other, optional uses of the digital signatures between partners.

If a trusted, self-signed certificate authority (CA) does not issue a digital-signing certificate, then the signing partner must send the public-key certificate out-of-band, such as through email, to the partner. The partner must import the certificate into PingFederate when configuring a connection to the signing partner for SSO, single log-out (SLO), or STS.

If a trusted CA signs the certificate and the signing partner chooses to embed the certificate in all signed messages, then the verifying partner uses the embedded certificate for signature verification, after validating it against the Subject Distinguished Name (DN) of the original certificate. Because validation only requires the Subject DN, the public-key certificate might not be sent out-of-band.



During the signature verification configuration, PingFederate extracts the Subject DN from the certificate when available.

The next section provides more information about the two alternative signature-verification trust models described above from the standpoint of the verifying partner.

Trust models

For validating digital signatures, PingFederate provides a selection of trust models in the administrative console for each partner connection, based on the certificate categories listed below. For each trust model, PingFederate always verifies the currency of the certificate and the validity of the message in the certificate specified. Additional checks depend upon the trust model selected.

Anchored certificate

In this case, certificates used for signature verification must be issued by a trusted CA, and the certificate chain must be verifiable recursively back to the root issuer. PingFederate validates the certificate, including recursive revocation checking (when enabled) back to the issuer, for all signed messages from the partner. By default, PingFederate also prompts for the Issuer DN of the certificates to mitigate potential man-in-the-middle attacks and to provide a means to isolate certificates used by different connections.

In addition, when choosing the anchored trust model, the incoming message must include the verification certificate for the signature. PingFederate uses that certificate to verify signatures from the partner if its Subject DN matches the partner's public certificate (as specified in the administrative console), the Issuer DN (if specified) matches one of the issuers in the chain, and the Issuer CA certificate is part of the trusted store. This feature provides a dynamic trust model that overcomes the problem of interrupting service to change out expired certificates.

Unanchored Certificate

In the unanchored certificate model, incoming signature verification exclusively uses the certificates imported for a connection into PingFederate or a secondary, backup certificate, either self-signed or trusted CA-issued, when specified. This verification does not apply to the certificate chain. However, when enabled, existing chains receive revocation checking as far as available.

Secure sockets layer

A certificate authority (CA)-signed SSL certificate identifies one or both ends of the federation. SSL/TLS provides an encrypted connection between the two parties to avoid exposing the content of a message. This promotes confidentiality and message integrity.

SAML SSL and TLS scenarios

SSL/TLS should be used in association with the SOAP responder URL and Single Sign-on Service located at an identity provider (IdP) site. On the service provider (SP) side, the Artifact Resolution Service should also use SSL/TLS. Optionally, SSL/TLS can also be used to secure communication between internal data stores and PingFederate and between the PingFederate security token service (STS) and web service client or provider applications.

The SSL/TLS server-client handshake involves negotiating cipher suites to use for encryption and decryption on each side of a secured transaction. You can find cipher suites in the following configuration files:

- com.pingidentity.crypto.SunJCEManager.xml
- com.pingidentity.crypto.AWSCloudHSMJCEManager.xml
- com.pingidentity.crypto.LunaJCEManager.xml
- com.pingidentity.crypto.NcipherJCEManager.xml
- com.pingidentity.crypto.BCFIPSJCEManager.xml

These cipher-suite configuration files are located in the <pf install>/server/default/data/ config-store directory. These files comment out weaker cipher suites. To ensure the most secure transactions, retain this cipher-suite configuration.



Important:

Due to the import restrictions of some countries, Oracle Server Java SE Runtime Environment (JRE) 8 has built-in restrictions on available cryptographic strength (key size). To use larger key sizes, enable the Java Cryptography Extension (JCE) unlimited strength jurisdiction policy. For more information, see the Java 8 release notes in Oracle's documentation.

For Oracle Java SE Development Kit 11, the JCE jurisdiction policy defaults to unlimited strength. For more information, see the Oracle JDK Migration Guide in Oracle's documentation.

Starting with version 9.1, PingFederate selects cipher suites based on the order that they appear in the cipher-suite configuration file for new installations. For upgrades, enable the same selection mechanism. For more information, see *Managing cipher suites* on page 820.

PingFederate browser-based single sign-on (SSO) uses three methods to authenticate connection partners making SOAP requests. For STS client SOAP authentication, configure a separate option using either or both of the first two methods listed here. Partners must agree upon the selection of methods and synchronize within IdP and SP federation implementations.

HTTP Basic authentication

Partners identify themselves by passing username and password credentials.

SSL client certificate authentication

Partners use SSL client certificates presented during SOAP request transactions. Each partner needs to import the other partner's certificate out-of-band For more information, see Manage SSL client keys and certificates on page 606.

Digital signatures

Partners sign the XML message transmitted through the SSL/TLS connection. The receiver verifies the signatures based upon the certificates configured for that connection. Each partner should import the others' certificates out-of-band. For more information, see Manage digital signing certificates and decryption keys on page 608.

Trusted certificates

PingFederate validates the trust of all certificates. PingFederate trusts certificates if the issuer's certificate is also in PingFederate's trusted certificate store. You must import the root certificate of the CA into PingFederate's trusted certificate store or into the Java runtime cacerts store.

Encryption

PingFederate supports the optional SAML 2.0 specification allowing for encryption of assertions, including security token service (STS) SAML tokens, which further enhances confidentiality when required.

For SAML 2.0 single sign-on (SSO) connections, you can choose to encrypt entire assertions or individual user attributes, including the user's name identifier. You can use signature verification and signing keys to encrypt and decrypt messages, respectively.

Hierarchical plugin configurations

PingFederate allows you to use a configuration of an adapter along with other PingFederate plugins as a parent instance for any created child instances.

After this, you can then modify the inherited configuration for child instances as needed. This feature provides easier management of adapter settings in cases where only small changes to an existing adapter or plugin configuration need to be made for a particular use case.

For example, different service provider (SP)-connection adapter instances might have their own identity provider (IdP) sign-on URLs, for branding or other application ownership reasons, while the majority of the other adapter configuration settings are the same. In this case, you might want to use a parent/child configuration to override the logon URLs.



Override adapter instances as part of mapping them into either SPor IdP connections for cases where overridden settings apply only to one particular connection configuration.

Any changes to a parent configuration propagate to its child or connection-based configurations so long as the derived instance has not overridden the changes.

In addition to adapters, PingFederate allows you to create parent/child configurations for the following plugin types:

- Access Token Management instances (see Access token management on page 549)
- Password Credential Validators (see Password Credential Validators on page 905)
- Identity Store Provisioners (see Configuring Identity Store Provisioners on page 653)

Identity mapping

Identity mapping is at the core of identity federation. One of the primary goals of SAML is to provide a way for an identity provider (IdP) to send a secure token, called the assertion, containing user-identity information that a service provider (SP) translates or maps to local user stores.

For browser-based single sign-on (SSO), PingFederate enables two modes of identity mapping between domains: account linking and account mapping.

For WS-Trust security token service (STS), PingFederate uses account mapping.

See subsequent topics for more information about these identity mapping options.

Account linking

Under the standards, use account linking for browser-based single sign-on (SSO) in cases where each domain maintains separate accounts for the same user.

Account linking uses the SAML assertion to create a persistent association between these distinct user accounts. The account link, or name identifier, such as an email address or identity provider (IdP)generated pseudonym, identifies individual users. When privacy is a concern, use pseudonyms because they prevent tracing back to a user's identity at the partner site.

During the user's first SSO request, the service provider (SP) prompts for local credentials, which enables the SP to link the name identifier contained within the assertion—either an open attribute or a pseudonym —with the user's local account. Subsequent SSO events will not prompt the user to authenticate with the SP because the SP federation server keeps a table associating remote users' name identifiers with local user accounts. The SP associates the link to the user's corresponding local account and provides access to the account without separate authentication.



PingFederate in the SP role uses a default, HSQLDB database to handle account linking. You can use your own datastore instead, as needed. For more information, see Account-linking datastores on page 901.



CAUTION:

Use the built-in HSQLDB only for trial or training environments. For testing and production environments, always use a secured external storage solution for proper functioning in a clustered environment.

Testing involving HSQLDB is not a valid test. In both testing and production, it might cause various problems due to its limitations and HSQLDB involved cases are not supported by Pingldentity.

The name identifier optionally includes additional attributes. When using a pseudonym as the account link, take care to send only general attributes, such as a user's organizational role or department, that will not compromise privacy.

Linking permission and defederation

The SAML specification also allows the SP application to build in user verification and approval of account linking and provides a means for the user to permanently cancel the linking, known as defederation. For more information, see /sp/defederate.ping on page 1142. A defederated user might later elect to reassociate with a local user account.

Under the SAML 2.0 specifications, an ildP configures PingFederate to enable a group of SPs, called an SP affiliation, to share the same persistent name identifier. For more information, see *SP affiliations* on page 471. An SP affiliation facilitates the use case where a number of business partners have an existing relationship and where sharing a single name identifier among all parties reduces the federation integration effort.

Account mapping

Account mapping, also called attribute mapping, enables a service provider (SP) to use PingFederate to perform a user lookup and map a user's identity dynamically based on one or more attributes received in the assertion.

Looking up the user always exposes the attributes. In other words, both the identity provider (IdP) and the SP know these attributes, such as an email address.

Account mapping achieves one-to-one mapping where individual user accounts exist on both sides of a federated connection or many-to-few mapping where IdP users without accounts at destination sites map to guest accounts or to a role-based general account.

For browser-based single sign-on (SSO), transient identifiers provide an additional level of privacy—virtual anonymity—by generating a different opaque ID each time the user initiates SSO. Transient IDs are often used in conjunction with federation role mapping to map the user to a guest account or to a role-based account based on the user's association with the IdP organization rather than personal attributes.

As with pseudonyms, additional attributes might be sent with the transient identifier. Again, take care to preserve privacy.

In B-to-B or B-to-E use cases where an administrator might create a user lookup on behalf of the user, the administrator might implement account mapping.

User attributes

Federation transactions require, at a minimum, the transmission of a unique piece of information, such as an email address, that identifies the user for identity mapping between security domains.

In addition to attributes used for identity mapping, the identity provider (IdP) can pass other user attributes in an assertion, including SAML tokens for web services. The service provider (SP) uses this supplemental information for several purposes. For example, the SP can use attributes to map and authorize the user into a specific role with associated site permissions or to customize the end application display for a more robust user experience.

The SP can also incorporate additional attributes prior to creating a session for the target application. This is common where the SP also maintains an account for the user and wants additional information for profiling or access-policy purposes.

Attributes must be carefully managed between IdPs and SPs. PingFederate facilitates the process by providing configuration steps that enable administrators to:

- Define and enforce attribute_contract for each partner connection.
- Define and retrieve attributes from the IdP adapter, authentication policy contracts, or security token service (STS) token processor to populate an attribute contract directly or use these attributes to look up additional attributes in IdP data stores.
- Define and enforce a set of required attributes needed by SP adapters or STS token generators to interface local systems or applications.
- Set up connections to local data stores.
- Configure specific attribute sources and lookups based on the data stores and map attributes into IdP assertions or into SP adapters or token generators used to interface target applications.
- Selectively mask attribute values recorded in transaction logs.

An attribute contract represents an agreement between partners about user attributes sent in a SAML assertion, a JSON web token (JWT), or an OpenID Connect ID token.

The contract is a list of case-sensitive attribute names. Partners must configure attribute contracts to match.



When privacy is required for sensitive attributes, you can configure PingFederate to mask their values in log files. For more information, see Attribute masking on page 95.

For an identity provider (IdP) or an OpenID Provider (OP), the attribute contract defines which attributes PingFederate sends in an assertion, a JWT, or an ID token. While all users authenticate to the partner through this fixed contract, the values used to fulfill the contract might differ from one user to the next. Relying on a combination of different data sources might also fulfill the attribute contract:

- The IdP adapter or security token service (STS) token processor
- An IdP attribute source, which identifies the location of individual attributes in a datastore
- Static text values for some attributes, or text values combined with variables
- Expressions (see Attribute mapping expressions on page 214)

For a service provider (SP) or an OpenID Connect Relying Party, the attribute contract defines the attributes PingFederate expects in a SAML assertion, an ID token, or from the UserInfo endpoint at the OP. To pass these attributes to the SP adapter or, for web services, to the SP token generator, configure PingFederate accordingly. For more information, see *Managing SP adapters* on page 649 or *Managing* token generators on page 975. In addition, you can configure PingFederate to use attributes to look up additional attributes in local data stores, which often help start a user session or create a local security token for web services. For more information, see Adapter contracts on page 94 or STS token contracts on page 94.

The attribute contract always contains the user identifier SAML SUBJECT in a SAML assertion and sub in a JWT or an ID token unless you are using account linking for browser-based single sign-on (SSO). This attribute is automatically included when creating a new contract.



Note:

You create attribute contracts on a per-connection basis. For example, if an SP has deployed two sessioncreation adapters for two separate applications, the IdP connection partner creates a single attribute contract. This single contract supplies all the attributes required by both SP adapters.

Name formats

By agreement with an SP partner, an IdP might specify a format, such as email, associated with the SAML SUBJECT. The SP might further require this information to facilitate handling of the format.

The partner agreement might also include a requirement for the IdP to provide format specifications associated with other attributes.

PingFederate provides a means for an IdP administrator to select from among standard subject, attribute formats, or both, depending on the relevant SAML specifications. An administrator also defines a customized selection of additional attribute formats. For more information, see Setting up an attribute contract on page 414.



Note:

For the WS-Trust IdP configuration, attribute-name formats remain unspecified. If needed however, an administrator might user a special variable in the attribute contract to set the subject-name format. For more information, see *Defining an attribute contract for IdP STS* on page 963. Browser-based SSO attribute contracts also use the same variable, but the feature has deprecated.

STS namespaces

By agreement with an SP partner for a WS-Trust STS connection, an IdP specifies an XML namespace to associate with an attribute, for example, to use claims-based authorization with WIF clients. For more information, see *WSC and WSP support* on page 73. The only attributes that allow specified namespaces belong to a WS-Trust IdP configuration using SAML 1.1 or SAML 1.1 for Office 365 as the default token type. For more information, see *Defining an attribute contract for IdP STS* on page 963.

Adapter contracts

An adapter contract represents an agreement between the PingFederate server and an external application.

In concert with the attribute contract between partners, adapter contracts specify the transfer of attributes. Adapter contracts consist of a list of case-sensitive attribute names.

On the identity provider (IdP) side of a federation, an IdP adapter supplies attributes to PingFederate for more information, see SSO integration kits and adapters on page 83 and Managing IdP adapters on page 383.

On the service provider (SP) side, adapters require adapter contract attributes to start a session with an application. Each security domain requires at least one adapter type. Then, you must configure an adapter instance for each target application. For more information, see *Managing SP adapters* on page 649.

Attributes from the attribute contract fulfill the adapter contracts on the SP side, possibly enhanced with other attributes from local data stores. For example, if the same security context controls several target applications and provides the same set of attributes to start a session for the user, you would deploy an adapter type and configure an adapter instance for each protected application. For more information, see *Managing target session mappings* on page 676.

Extended adapter contract

When PingFederate deploys an adapter type, it creates adapter contracts. Developing these adapters "hard-wires" them to look up or set a specific set of attributes. Attribute requirements might change after deployment. To streamline adjustment of adapter contracts, PingFederate allows an administrator to add additional attributes to the adapter instance through the administrative console, called extended adapter contracts.

STS token contracts

Similar to an adapter contract for broswer-based single sign-on (SSO), A security token service (STS) token-processor or token-generator contract represents an agreement between the PingFederate server and an external application in the context of a web services transaction.

In concert with the attribute contract between partners, token contracts specify the transfer of attributes, consisting of a list of case-sensitive attribute names.

On the identity provider (IdP) side of a federation, PingFederate receives token-processor attributes. For more information, see *Token processors and generators* on page 72 and *Managing token processors* on page 953.

On the service provider (SP) side, a token generator requires token-generator contract attributes to pass identify information from the token to the web service client application. Each security domain requires

at least one token generator type. Then a token-generator instance must be configured for each target application. For more information, see *Managing token generators* on page 975. If several target applications are controlled by the same security context and can receive the same set of attributes for the user, you would deploy a token generator type and configure a token generator instance for each target application. For more information, see *Managing SP token generator mappings* on page 979.

Extended token generator contract

When PingFederate deploys a token-generator type, it creates token-generator contracts. When developed, these token generators are "hard-wired" to look up or set a specific set of attributes. After deployment, your attribute requirements might change. To streamline adjustment of token-generator contracts, PingFederate allows an administrator to add additional attributes to the token-generator instance through the administrative console. These adjustments are called extended token-generator contracts.

Datastores

Datastores represent external systems that store user attributes and other data.

Once defined, PingFederate configurations retrieve user attributes from datastores for contract fulfillment and token authorization in various use cases. PingFederate configurations write certain records or log messages to datastores. PingFederate supports a wide variety of database servers and directory servers. As needed, the PingFederate SDK supports the creation of custom drivers for connecting to other types of data repositories, such as flat files or SOAP-connected databases. For more information, see the Javadoc for the CustomDataSourceDriver interface, the SamplePropertiesDataStore.java file for a sample implementation, and the SDK Developer's Guide on page 1021 for build and deployment information.



The Javadoc for PingFederate and the sample implementation are in the <pf install>/ pingfederate/sdk directory.

Attribute masking

At runtime PingFederate logs user attributes. To preserve user privacy, you can mask the values of logged attributes.

For more information about log files, see PingFederate log files on page 761. PingFederate provides this masking capability at all points where the server logs attributes. These points include:

- Datastore lookup at either the identity provider (IdP) or service provider (SP) site. For more information, see *Datastores* on page 867.
- Retrieval of attributes from an IdP adapter or token processor. For more information, see Setting pseudonym and masking options on page 385 and Setting attribute masking on page 958.
- SP-server processing of incoming attributes based on the single sign-on (SSO) attribute contract. For more information, see *Defining an attribute contract* on page 674.



Note:

The SAML Subject ID is not masked; the SAML specifications provide for either pseudonymous account linking or transient identification to support privacy for the Subject ID. For more information, see Account linking on page 91.

For information about XASP, see Attribute Query and XASP on page 43.



Important:

Many adapter implementations, along with other product extensions, can independently write unmasked attribute values to the PingFederate server log. PingFederate does not control these implementations. If using such a component raises a concern about sensitive attribute values, you can adjust the component's logging threshold in log4j2.xml to prevent the recording of attributes.

About token authorization

PingFederate provides an optional configuration known as token authorization to evaluate user attributes as well as other runtime variables, such as authentication context, for authorization purposes.

Token authorization provides a way for administrators to extend access policy directly to many areas, such as browser single sign-on (SSO), security token service (STS), and OAuth events, by conditionally allowing or disallowing the issuance of relevant security tokens such as SAML assertions, STS tokens, OAuth access tokens, or session cookies. The option is also available for extending authorization policy to attribute-query responses, identity provider (IdP) adapter contracts, and authentication policy contracts.

Administrators can configure token authorization using Issuance Criteria windows immediately following the configuration of attribute mapping at all applicable points in the administrative console. See Defining issuance criteria for IdP Browser SSO on page 423 as an example.

Issuance criteria

The token-authorization configuration consists of rules that evaluate attribute values against selected conditions. Depending on the type of configuration that contains the token-authorization setup, choose from among several sources for the attributes. The sources always consist of all of those available for attribute mapping, including configured data stores and runtime information related to the context of an event. In addition, the sources use values of mapped attributes to provide access to any plain-text mappings or the runtime results of any attribute mapping expressions.



When more than one condition is configured, all are evaluated and all the conditions must be met at runtime (evaluated as true) for authorization to succeed and processing to continue. In cases where you might need "or" conditions or layered evaluations, you can create one or more attribute mapping expressions.



Note:

When authorization fails and a transaction halts, the system passes back a configurable Error Result code, potentially to an application layer or as a variable on a PingFederate user-facing template. How this code is interpreted depends on the use case and application integration.

Single and multivalued conditions

Each token-authorization configuration provides a choice of conditions for evaluating attribute values:

- equal to
- equal to (case insensitive)

- equal to DN
- not equal to
- not equal to (case insensitive)
- not equal to DN
- multi-value contains
- multi-value contains (case insensitive)
- multi-value contains DN
- multi-value does not contain
- multi-value does not contain (case insensitive)
- multi-value does not contain DN



Note:

The first six conditions are intended for single-value attributes. Use one of the multi-value ... conditions when you want PingFederate to validate whether one of the attribute values matches or does not match the specified value. Using a single-value condition when an attribute has multiple values causes the criteria to fail consistently.

User provisioning

PingFederate provides cross-domain user provisioning and account management.

User provisioning is an important aspect of identity federation. When organizations enable SSO for their users, they must ensure that some form of account synchronization is in place. Automated user provisioning features within PingFederate free administrators from having to devise a manual strategy for this.

Provisioning support takes different forms, depending on what role PingFederate plays in an identity federation, and you can configure it either in conjunction with partner SSO connections or separately:

 IdP sites support automatic provisioning and maintaining user accounts at service provider (SP) sites that have implemented the system for cross-domain identity management (SCIM) or at selected software as a service (SaaS) providers. For more information, see the next section, Outbound provisioning for IdPs on page 97.

For information about SCIM, see www.simplecloud.info.

 When PingFederate is configured as a SP, it supports provisioning and managing user accounts and groups for your own organization automatically by using the standard SCIM protocol or by using identity information received during SSO events from SAML assertions. For more information, see Provisioning for SPs on page 98.

Outbound provisioning for IdPs

For identity provider (IdP) sites, PingFederate provides built-in automated provisioning and user-account management to system for cross-domain identity management (SCIM)-enabled services providers and to selected software as a service (SaaS) providers through their proprietary provisioning APIs.

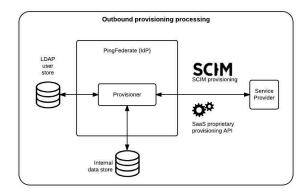
User provisioning is an important aspect of identity federation. When organizations enable SSO for their users, they must ensure that some form of account synchronization is in place. Automated user provisioning features within PingFederate free administrators from having to devise a manual strategy for this.

Outbound provisioning also provides an automated means of account disabling or deprovisioning, which might be of key importance to system auditors.



Z Tip:

With outbound provisioning enabled, the PingFederate runtime engine, the provisioner, polls the IdP organization's user store periodically. The server uses a separate database to monitor the state of the user store and keeps user data synchronized between the organization and the target service provider, as illustrated in the following diagram.



LDAP user store

PingFederate provides built-in support for PingDirectory, Microsoft Active Directory, Oracle Unified Directory, and Oracle Directory Server; pre-configuration of many provisioning settings uses templates. Although Ping Identity has only formally tested these datastores for support, other LDAP datastores will likely work as well.

Internal datastore

PingFederate is tested with Amazon Aurora (MySQL and PostgreSQL), Microsoft SQL Server, Oracle Database, Oracle MySQL, and PostgreSQL as internal provisioning datastores. A demonstration-only, embedded HSQLDB database is installed by default. Scripts to aid setup are in the directory <pf install>/pingfederate/server/default/conf/provisioner/sqlscripts.



CAUTION:

Use the built-in HSQLDB only for trial or training environments. For testing and production environments, always use a secured external storage solution for proper functioning in a clustered environment.

Testing involving HSQLDB is not a valid test. In both testing and production, it might cause various problems due to its limitations and HSQLDB involved cases are not supported by Pingldentity.

Provisioning for SPs

PingFederate offers two provisioning options when configured as a service provider (SP).

User provisioning is an important aspect of identity federation. When organizations enable SSO for their users, they must ensure that some form of account synchronization is in place. Automated user provisioning features within PingFederate free administrators from having to devise a manual strategy for this.

When configured as an SP, PingFederate offers two provisioning options:

Inbound provisioning

SCIM inbound provisioning provides support for incoming SCIM messages containing requests to create, read, update, delete, or deactivate user and group records in Microsoft Active Directory datastores or

custom user stores through the Identity Store Provisioners. PingFederate supports SCIM attributes in the core schema and custom attributes through a schema extension. Configuring this provisioning feature has two options: by itself or in conjunction with SSO or other connection types.

In effect, inbound provisioning provides an organization with a dedicated SCIM service provider, which routes user-managment requests to an organization's centralized user store. The requests usually originate from trusted applications within an organization, such as a human-resources on-boarding software as a service (SaaS) product, or from trusted partner identity providers (IdPs).

For setup information, see Configuring SCIM inbound provisioning on page 709. To integrate inbound provisioning with custom user stores, see Configuring Identity Store Provisioners on page 653. For application-development information about using PingFederate endpoints for SCIM provisioning, see SCIM inbound provisioning endpoints on page 1144.

Just-in-time provisioning

At an SP site, PingFederate creates and updates local user accounts in an external LDAP directory or Microsoft SQL Server as part of SSO processing, called Just-in-time (JIT) provisioning or, formerly, Express Provisioning. When provisioning requires local accounts, this feature allows SPs to maintain accounts for users who authenticate through IdP partners without having to provision accounts manually.

When configured, the PingFederate SP server writes user information to the local user store using attributes from the incoming SAML assertion. For SAML 2.0 partner connections, supplement assertion attributes with user attributes returned from an Attribute Query.

PingFederate also updates existing user accounts based on assertions. Using this option, PingFederate adds or overwrites attributes for a local user account each time PingFederate processes SSO for a user.

For information about enabling JIT Provisioning, see Choosing IdP connection options on page 665. For configuration information, see Configuring just-in-time provisioning on page 699.

Customer identity and access management

PingFederate empowers administrators to deliver a secure and easy-to-use customer authentication, registration, and profile management solution.

This solution leverages the HTML Form Adapter to offer users the options to authenticate through thirdparty identity providers, self-register as part of the sign-on experience, and manage their accounts through a self-service profile management page. The registration and profile management pages are fully customizable and localizable, which allows administrators to present a consistent branding experience based on the needs of the users and the organizations.

Federation hub use cases

Configuring PingFederate as a federation hub accomplishes two primary functions.

Configure PingFederate as a federation hub to:

- Bridge partners using different federation protocols to circumvent partner or application limitations.
- Multiplex a connection for multiple partners to reduce costs and expand use cases.

As a federation hub, PingFederate bridges browser-based single sign-on (SSO) between identity providers and service providers. It stands in the middle of the SSO and single log-out (SLO) flow, acting as the service provider (SP) for the identity providers and as the identity provider (IdP) for the service providers. The four use cases are:

- Bridging an IdP to an SP
- Bridging an IdP to multiple SPs
- Bridging multiple IdPs to an SP
- Bridging multiple IdPs to multiple SPs

PingFederate also supports protocol translation among SAML 1.0, 1.1, 2.0, OpenID Connect, and WS-Federation. For SAML-based connections, this also means it is possible to bridge between various bindings between identity providers and service providers.

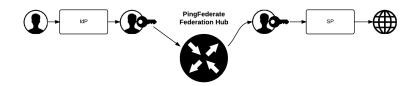
The federation hub capability deploys alongside with other OAuth use cases, IdP connections, SP connections, or any combination of them, to your partners. This flexibility helps in streamlining your federation infrastructure and reducing operating costs.

Bridging an IdP to an SP

PingFederate bridges single sign-on (SSO) and single log-out (SLO) transactions between an identity provider (IdP) and a service provider (SP).

About this task

If you have a legacy IdP system only capable of sending SAML 1.1 assertions through POST and an SP that requires SAML 2.0 assertions through the artifact binding, configuring the federation hub allows PingFederate to consume inbound SAML 1.1 assertions by POST, translate them to SAML 2.0 assertions, and send them through the artifact binding the SP.



Steps

- 1. Create a contract to bridge the attributes between the IdP and the SP. For more information, see *Federation hub and authentication policy contracts* on page 103.
- Create an IdP connection between the IdP and PingFederate, the federation hub as the SP, and add the applicable authentication policy contracts to the IdP connection on the **Target Session Mapping** tab
- Create an SP connection between PingFederate, the federation hub as the IdP, and the SP and add to the SP connection the corresponding authentication policy contract on the **Authentication Source Mapping** window.
- 4. Work with the IdP to connect to PingFederate, the federation hub, as the SP.
- 5. Work with the SP to connect to PingFederate, the federation hub, as the IdP.

Bridging an IdP to multiple SPs

PingFederate bridges single sign-on (SSO) and single log-out (SLO) transactions between an identity provider (IdP) and multiple service providers (SPs).

About this task

For example, your company wants to route federation requests from a recently acquired subsidiary through its federation infrastructure. PingFederate multiplexes one IdP connection to multiple SP connections to the desired SPs. The federation hub consumes assertions from the subsidiary and creates new assertions to the respective SPs.

Steps

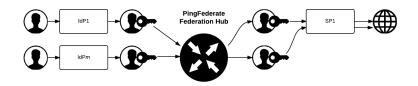
- 1. For each SP, create a contract to the IdP. For more information, see *Federation hub and authentication policy contracts* on page 103. Because each SP likely requires a unique set of attributes, you will need to create multiple contracts.
- 2. Create an IdP connection between the IdP and PingFederate, the federation hub as the SP.
- 3. Add the applicable authentication policy contract(s) to the IdP connection on the **Target Session Mapping** window.
- 4. For each SP, create an SP connection between PingFederate, the federation hub as the IdP, and the SP.
- 5. Add the corresponding authentication policy contract to the SP connection on the **Authentication Source Mapping** window.
- 6. For each SP supporting the SAML IdP-initiated SSO profile, map the expected target resources to the corresponding SP connections on the **Applications # Integration # Target URL Mapping** window.
- 7. Work with the IdP to connect to PingFederate, the federation hub as the SP.
- 8. Work with each SP to connect to PingFederate, the federation hub as the IdP.

Bridging multiple IdPs to an SP

With PingFederate, you can bridge single sign-on (SSO) and single log-out (SLO) transactions between multiple identity providers (IdPs) and a service provider (SP).

About this task

For example, you are tasked to provide federated access to resources on Microsoft SharePoint for various business partners. With PingFederate, you can multiplex one SP connection, to SharePoint, to multiple IdP connections for all your business partners. The federation hub can also, as needed, translate SAML assertions from the business partners to WS-Federation security tokens and send them over to SharePoint.



Steps

- 1. Create a contract to bridge the attributes between the IdPs and the SP. For more information, see *Federation hub and authentication policy contracts* on page 103.
 - You likely need only one contract unless the SP requires a different set of attributes from each IdP.
- 2. For each IdP, create an IdP connection between the IdP and PingFederate, the federation hub as the SP.

- 3. On the Target Session Mapping window, add the applicable authentication policy contracts to the IdP connection.
- 4. On the **Selectors** window, configure an authentication selector. For example, see an instance of the Identifier First Adapter on page 305 to map each IdP to the corresponding IdP connection in an authentication policy.
- 5. Create an SP connection between PingFederate, the federation hub as the IdP, and the SP.
- 6. Add the corresponding authentication policy contract to the SP connection on the **Authentication** Source Mapping window.



Important:

PingFederate includes the Entity ID of the original IdP (Authenticating Authority) in SAML 2.0 assertions so that the SP can determine the original issuer of the assertions. This is especially important when bridging multiple IdPs to one SP—the SP should take the information about the original issuer into consideration before granting access to protected resources.

For SAML 1.x assertions and WS-Federation security tokens, you can add an attribute on the Attribute Contract window and then map Context: Authenticating Authority as the attribute value on the Attribute Contract Fulfillment window.

For information about Authenticating Authority, see section 2.7.2.2 Element < AuthnContext> in the SAML 2.0 specification.



If the SP does not take action based on Authenticating Authority, depending on the attributes from the IdPs, you can define validation rules on the **Issuance Criteria** window to protect against user impersonation between IdPs.

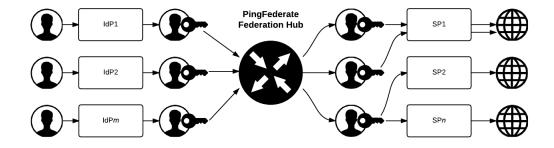
- 7. Work with each IdP to connect to PingFederate, the federation hub as the SP.
- 8. Work with the SP to connect to PingFederate, the federation hub as the IdP.

Bridging multiple IdPs to multiple SPs

PingFederate can bridge single sign-on (SSO) and single log-out (SLO) transactions between multiple identity providers (IdPs) and service providers (SPs).

About this task

This PingFederate federation hub use case is a combination of Bridging an IdP to multiple SPs on page 100 and Bridging multiple IdPs to an SP on page 101.



Steps

- 1. Create multiple contracts to bridge the attributes between the IdPs and the SPs. For more information, see Federation hub and authentication policy contracts on page 103.
- 2. For each identity provider, create an IdP connection between the IdP and PingFederate, the federation hub as the SP.
- 3. Add the applicable authentication policy contracts to the IdP connection in the Target Session Mapping window.
- 4. In the **Selectors** window, configure an authentication selector to map each IdP to the corresponding IdP connection in an authentication policy. For example, see an instance of the *Identifier First Adapter*
- 5. For each service provider, create an SP connection between PingFederate, the federation hub as the IdP. and the SP.
- 6. Add the corresponding authentication policy contract to the SP connection in the **Authentication** Source Mapping window.



Important:

PingFederate includes the Entity ID of the original IdP, Authenticating Authority, in SAML 2.0 assertions so that the SP can determine the original issuer of the assertions. This is especially important when bridging multiple IdPs to one SP—the SP should take the information about the original issuer into consideration before granting access to protected resources.

In the Attribute Contract window, add an attribute for SAML 1.x assertions and WS-Federation security tokens. Then, in the Attribute Contract Fulfillment, map Context: Authenticating **Authority** as the attribute value.

For information about Authenticating Authority, see section 2.7.2.2 Element < AuthnContext> in the SAML 2.0 specification.



Note:

If the SP does not take action based on Authenticating Authority, in the Issuance Criteria window, you can define validation rules to protect against user impersonation between IdPs.

- 7. For each SP supporting the SAML IdP-initiated SSO profile, map the expected target resources to the corresponding SP connections in the Applications # Integration # Target URL Mapping window.
- 8. Work with each IdP to connect to the federation hub as the SP.
- 9. Work with each SP to connect to the federation hub as the IdP.

Federation hub and authentication policy contracts

PingFederate uses two connections to bridge an identity provider (IdP) to a service provider (SP). It fuses these two connections together by using an authentication policy contract, formerly known as a connection mapping contract, as the medium to carry user attributes from the IdP to the SP.

The two connections are:

- An IdP connection where end users authenticate and PingFederate, the federation hub, is the SP
- An SP connection to the target application where PingFederate, the federation hub, is the IdP

Each authentication policy contract comes with one default attribute, subject. You can extend the contract to include additional attributes as needed. In most federation hub use cases, you configure PingFederate to pull attribute values from inbound assertions into the authentication policy contract in an IdP connection and to push those values from the authentication policy contract into the outbound assertions through an SP connection. For advanced use cases, you can configure the IdP connections, SP connections, or both, to look up values from multiple data store instances.

When bridging one IdP to one SP, you must create one authentication policy contract and associate the contract with both the IdP connection and the SP connection.

When bridging one IdP to multiple SPs, you need to create an authentication policy contract per SP because each SP likely requires a different set of attributes. In addition, you need to map all the authentication policy contracts into the IdP connection and add the respective authentication policy contracts to each SP connection to the SP.

When bridging multiple IdPs to one SP, you likely need only one contract unless the SP requires a different set of attributes from each IdP. In addition, you need to add the authentication policy contract to the SP connection and the applicable IdP connections.

You can manage authentication policy contracts on the **Policy Contracts** window (**Authentication** # **Policies** # **Policy Contracts**).

Federation hub and virtual server IDs

PingFederate manages the federation hub differently based on how the server provider (SP) connection uses virtual server IDs.

PingFederate uses two connections to bridge an identity provider (IdP) to a SP:

- An IdP connection where end users authenticate and PingFederate, the federation hub, is the SP
- An SP connection to the target application where PingFederate, the federation hub, is the IdP

Generally speaking, PingFederate consumes assertions from the IdP through the IdP connection and generates new assertions to the SP through the SP connection.

If the SP connection does not use a virtual server ID, the issuer of the assertions to the SP is the ID defined for the protocol between PingFederate, the federation hub as the IdP, and the SP.

If the SP connection uses multiple virtual server IDs for the purpose of connecting to multiple environments serviced by the same partner using one connection, PingFederate automatically retains information about AuthnRequest messages sent to the virtual server ID specific endpoint for SP-initiated single sign-on (SSO). When the IdP returns the corresponding assertions to PingFederate as the SP, PingFederate retrieves the preserved information and uses that specific virtual server ID as the issuer in assertions sent to the SP. For IdP-initiated SSO, the issuer of the assertions to the SP is the default virtual server ID.

Federation planning checklist

An essential first step in establishing an identity federation involves discussions and agreements between you and your connection partners. The sections below comprise a partial checklist of items that should be coordinated before you deploy PingFederate.

Standards and specifications

Choose which federation protocols your deployment will support. For SAML single sign-on (SSO) configurations, decide which profiles and bindings will be used. For more information, see *Supported Standards*.

Signing and validation

Decide which SAML messages—assertions, responses, requests—will be digitally signed and how the messages will be verified by your federation partner. If messages are signed, decide how certificates will be exchanged, for example, secure email. For more information, see *Security infrastructure* on page 86.

Back-channel security

Determine what type of SOAP channel authentication will be used: Basic or SSL/TLS. If SSL/TLS is used, determine whether server-only or both server and client certificates will be needed and how they will be managed. Also decide what level of security will be required for connections to back-end datastores or identity management systems.

Trusted certificate management

Determine whether both partners are using SSL/TLS runtime certificates, signing certificates, or both that have been signed by a major certificate authority (CA). If self-signed certificates or nonstandard CAs are used, the signed certificates must be exchanged and imported into Trusted Certificate stores. Also, determine whether you want to adopt a trust model that uses embedded certificates. For more information, see Digital signing policy coordination on page 88.

Deployment

Decide how PingFederate fits into your existing network. Also, determine whether high-availability, failover options, or both, are required. For more information, see the PingFederate Server Clustering Guide on page 182.

Federation server identification

Determine how you and your partners will identify your respective federation deployments. Under federation standards, both the sender, the identity provider (IdP), and the receiver, service provider (SP), of an assertion must be uniquely identified within the identity federation. For more information, see Configuration data exchange on page 107.

With PingFederate, you define a unique ID for each supported protocol. For more information, see Specifying federation information on page 838. Optionally, you can also use a list of multiple virtual server IDs on a connection-by-connection basis. For more information, see =.



PingFederate also provides for virtual host names, which differ from virtual server IDs but are not mutually exclusive; they are intended for use when your network configuration is such that you receive federation messages under more than one domain name. For more information, see Configuring virtual host names on page 855.

Server clock synchronization

Ensure that both the SP and IdP server clocks are synchronized. SAML messages and security token service (STS) tokens provide a time window that allows for small synchronization differentials. However, wide disparities will result in assertion or request time-outs.

User data stores

Identify the type of datastore that contains user data when needed. For more information, see *Datastores* on page 95.

Web application and session integration

Decide how PingFederate as an IdP receives subject identity information, either from an STS token or a user session.

For an SP, decide how PingFederate will forward user identity information to the destination web application or system to start a session. For more information, see SSO integration kits and adapters on page 83 and Token processors and generators on page 72.

Transaction logging

PingFederate provides basic transaction logging and monitoring. Decide whether transaction logging should be integrated with a systems management application and whether you have regulatory compliance requirements that affect your logging processes. For more information, see PingFederate log files on page 761.

Identity mapping

For browser-based SSO, decide whether you will use PingFederate to link accounts on your respective systems using a persistent name identifier, or whether you will use account mapping. For more information, see *Identity mapping* on page 91.

Attribute contract agreement

If your federation partnership will not use account linking, or will not use it exclusively, then you and your partner must agree on a set of attributes that the IdP will send in an assertion for either SSO or web service access. For more information, see *Attribute contracts* on page 93

Metadata exchange

If you are using SAML, decide whether you will use the metadata standard to exchange XML files containing configuration information. PingFederate makes it easy to use this protocol, which provides a significant shortcut to setting up your partner connections.

Multiple virtual server IDs

Virtual server IDs provide more configuration flexibility in cases where you need to identify your server differently when connecting to a partner in one connection for multiple environments or in multiple connections where the partner also supports multiple federation IDs.

Connecting to a partner in one connection

This is a use case where you need to connect to multiple environments serviced by the same partner using one federation ID, multiplexing one service provider (SP) connection to access multiple subdomain accounts in Microsoft Office 365.

Suppose both the marketing and the engineering departments of contoso.com, the identity provider (IdP), have their own departmental subdomains, marketing.contoso.com and engineering.contoso.com. They are both registered in Office 365, the SP, under the parent domain, contoso.com.

To include both marketing.contoso.com and engineering.contoso.com as the virtual server IDs in the Office 365 SP connection, configure the PingFederate IdP server. Each virtual server ID has its own set of protocol endpoints obtained in the connection metadata. For more information, see *Metadata export* on page 860 and *System-services endpoints* on page 1151.

After providing the protocol endpoints information to Office 365, when Office 365 sends login requests to PingFederate, PingFederate picks the correct IdP adapter to authenticate the end users based on the virtual server ID in the requests.

For each successful login, PingFederate builds an assertion with the issuer set to the corresponding virtual server ID. When Office 365 receives the assertion, it creates the end user session with the right subdomain settings based on the issuer value in the assertion.

Connecting to a partner in multiple connections

In this use case, you connect to your partner in multiple connections. In each connection, you identify yourself and your partner differently.

For example, you as the SP provide separate environments for the end users based on their regions. Your IdP operates in two regions, Europe (EU) and North America (NA). Their federation IDs are eu.idp.local and na.idp.local respectively.

In the PingFederate SP server, you can create two IdP connections to federate identities for end users from both regions as follows.

	Partner's federation ID	Your virtual server ID
IdP connection #1	eu.idp.local	idp-eu.sp.tld

	Partner's federation ID	Your virtual server ID
IdP connection #2	na.idp.local	idp-na.sp.tld

Based on the issuer (the partner's federation ID) and the audience values (your virtual server ID), PingFederate determines at runtime which IdP connection the assertion is intended for, validates as per the connection settings, and passes attribute values to the SP adapter to create the end-user session.

Working with multiple virtual server IDs

You can assign virtual server IDs either as an IdP during configuration of an SP connection or as an SP configuring an IdP connection for both Browser single sign-on (SSO) Profiles and WS-Trust security token service (STS) for access to identity-enabled web services. For more information, see *Identifying the SP* on page 407 and *Identifying the partner* on page 668.

If a connection has only one virtual server ID, it becomes the default virtual server ID for the connection. If the list contains several entries, you must specify one of them as the default virtual server ID for that connection. The connection uses the default virtual server ID when a request does not include virtual server ID information. For more information, see IdP endpoints on page 1133 for an IdP or SP endpoints on page 1138 for an SP.

In a connection with multiple virtual server IDs, you can restrict each adapter added to the connection to certain virtual server IDs to enhance the end-user experience. For more information, see Restricting an authentication source to certain virtual server IDs on page 420 and Restricting a target session to certain virtual server IDs on page 678.



Restrict each token processor or token generator added to a WS-Trust STS SP connection or IdP connection. For more information, see Restricting a token processor to certain virtual server IDs on page 967 or Restricting a token generator to certain virtual server IDs on page 981.



Important:

To protect against unauthorized access, configure Issuance Criteria to verify virtual server ID in conjunction with other conditions, such as group membership information. For more information, see *Defining issuance* criteria for IdP Browser SSO on page 423 or Defining issuance criteria for SP Browser SSO on page 681.

Configuration data exchange

If your partner's deployment does not produce or consume a metadata file that conforms to SAML metadata specifications, you might need to exchange connection information manually. If the deployment does not use metadata, some common configuration details must be exchanged.

Identity provider (IdP) to service provider (SP)

If you are the IdP, your SP partner will need some or all of the following connection information, depending upon which profiles and bindings you configure:

- Unique ID—Identifies the IdP that issues an assertion or other SAML message. For SAML 2.0, the ID is the IdP entity ID; for SAML 1.x, it is the IdP issuer; for WS-Federation, it is the IdP realm.
 - PingFederate also supports the optional use of virtual IDs. For more information, see Federation Server Identification.
- SOAP artifact resolution URL—The endpoint your site uses to receive an SP's SOAP requests when the artifact binding is used.

- Single logout (SLO) service URL—The destination of SLO request messages.
- Single sign-on (SSO) service URL—The endpoint where you receive and process assertions.

SP to IdP

If you are the SP, your IdP partner will need some or all of the following connection information depending upon which profiles and bindings you configure:

• Unique ID—Identifies the SP. For SAML 2.0, the ID is the entity ID; for SAML 1.x, it is the SP's audience; for WS-Federation, it is the SP's realm.

PingFederate also supports the optional use of virtual IDs. For more information, see *Federation Server Identification*.

- SOAP artifact resolution service URL—The endpoint to use for SOAP requests when the artifact binding is used.
- Single logout service URL (SAML 2.0)—The destination of SLO request messages.
- Assertion consumer service URL—The location where the SP receives assertions.
- Target URLs—The URLs for the protected resources that a user is trying to access.

Mutual settings between parties

The parties must mutually determine the settings. These settings might include:

- Attributes—User information sent in an assertion. For more information, see *User attributes* on page 92.
- Signing certificates—SAML and WS-Federation protocols specify a number of conditions built into the PingFederate connection-setup windows that might or might not require digital signatures.
- SOAP connection type and authentication style—For SAML connections using the back channel, such
 as artifact binding, HTTP Basic authentication, SSL client certificate authentication, digital signatures,
 or some combination of the three is required. You and your partner must exchange the necessary
 credentials, certificates, and signing keys.

Installing PingFederate

PingFederate operates as a standalone server based on Java EE application server technology. This section shows you how to properly install PingFederate.

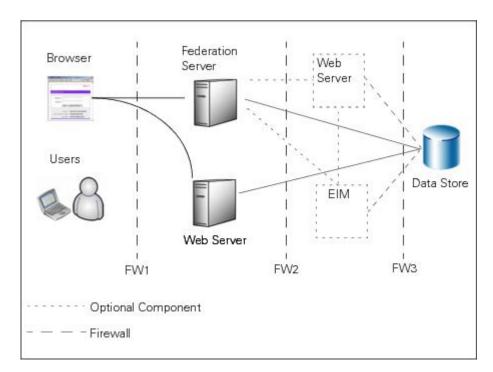
A new installation involves the following tasks:

- Determining the deployment architecture.
- Reviewing system and port requirements.
- Installing a Java runtime environment.
- Installing PingFederate.
- Completing the Initial Setup wizard.

Deployment options

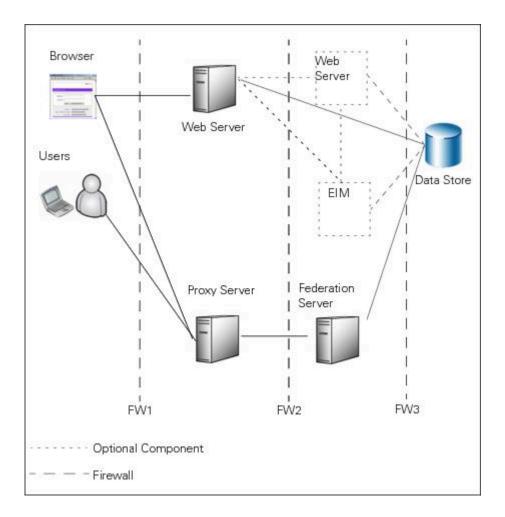
Depending on your needs and infrastructure capabilities, you can choose a standalone or proxy configuration. For information about configuring proxy settings, see *Configure incoming proxy settings* on page 635 and *Configuring forward proxy server settings* on page 833.

The following diagram illustrates a PingFederate install in a demilitarized zone (DMZ):



In this configuration, the users access PingFederate through a web application server, an enterprise identity management (EIM) system, or both. PingFederate, in turn, retrieves information from a datastore to use in processing the transaction.

You can also deploy PingFederate with a proxy server. The following diagram depicts a proxy-server configuration in which users and web browsers access the proxy. The proxy, in turn, communicates with PingFederate to request SSO.



System requirements

PingFederate supports the following system requirements. This section lists recommended versions and requirements.

Operating systems and virtualization



Note:

PingFederate is tested with default configurations of operating-system components. If your organization customizes implementations or installs third-party plug-ins, deployment efforts might affect the PingFederate server.

Operating systems

- Amazon Linux 2
- Canonical Ubuntu 16.04 LTS
- Canonical Ubuntu 18.04 LTS
- Microsoft Windows Server 2016
- Microsoft Windows Server 2019
- Oracle Enterprise Linux 7.7 (Red Hat compatible kernel)
- Oracle Enterprise Linux 8.1 (Red Hat Compatible Kernel)
- Red Hat Enterprise Linux ES 7.7

- Red Hat Enterprise Linux ES 8.1
- SUSE Linux Enterprise 12 SP5
- SUSE Linux Enterprise 15 SP1

Docker support

Docker version: 19.03.8

Host operating system: Ubuntu 18.04 LTS

Kernel: 4.15.0-1063-aws

Virtualization

Although Ping Identity does not qualify or recommend any specific virtual-machine (VM) or container products other than those listed above, PingFederate has run well on several, including Hyper-V, VMWare, and Xen.



Note:

The list of products is provided for example purposes only. We view all products in this category equally. Ping Identity accepts no responsibility for the performance of any specific virtualization software and in no way guarantees the performance, interoperability, or both of any VM or container software with its products.

Java environment

- Amazon Corretto 11
- Amazon Corretto 8
- OpenJDK 11
- Oracle Java SE Development Kit 11 LTS
- Oracle Java SE Runtime Environment (Server JRE) 8



Ping Identity Java Support Policy applies. For more information, see Java Support Policy in the Ping Identity Knowledge Base.

Browsers

Runtime server

- Apple Safari
- Google Chrome
- Microsoft Edge
- Microsoft Internet Explorer 11
- Mozilla Firefox
- Apple iOS 13 (Safari)
- Google Android 10 (Chrome)

Administrative server

- Google Chrome
- Microsoft Internet Explorer 11 or later
- Mozilla Firefox

TLS protocol

Runtime server and administrative server

TLS 1.2 and 1.3



TLS 1.3 requires Java 11.

Datastore integration

User-attribute lookup

- PingDirectory 7.0, 7.2, 7.3, 8.0, 8.1
- Microsoft Active Directory 2016
- Oracle Directory Server Enterprise Edition 11g
- Oracle Unified Directory 12c
- Amazon Aurora (MySQL 5.6.10a)
- Amazon Aurora (PostgreSQL 10.7)
- Microsoft SQL Server 2016 SP2 and 2017
- Oracle Database 12c Release 1 (12.1.0.2.0)
- Oracle Database 19c
- Oracle MySQL 8.0
- PostgreSQL 9.6.15 and 11.5

SaaS or SCIM outbound provisioning

Provisioning channel data source

- PingDirectory 7.0, 7.2, 7.3, 8.0, 8.1
- Microsoft Active Directory 2016
- Oracle Directory Server Enterprise Edition 11g
- Oracle Unified Directory 12c

Provisioning internal datastore

- Amazon Aurora (MySQL 5.6.10a)
- Amazon Aurora (PostgreSQL 10.7)
- Microsoft SQL Server 2016 and 2017
- Oracle Database 12c Release 1
- Oracle Database 19c
- Oracle MySQL 8.0
- PostgreSQL 9.6.15 and 11.5

SCIM inbound provisioning

- Microsoft Active Directory 2016
- Custom implementation through the PingFederate SDK

Just-in-time (JIT) inbound provisioning

- PingDirectory 7.0, 7.2, 7.3, 8.0, 8.1
- Microsoft Active Directory 2016

- Oracle Directory Server Enterprise Edition 11g
- Oracle Unified Directory 12c
- Microsoft SQL Server 2016 SP2 and 2017

Account linking

- PingDirectory 7.0, 7.2, 7.3, 8.0, 8.1
- Microsoft Active Directory 2016
- Oracle Directory Server Enterprise Edition 11g
- Oracle Unified Directory 12c
- Amazon Aurora (MySQL 5.6.10a)
- Amazon Aurora (PostgreSQL 10.7)
- Microsoft SQL Server 2016 SP2 and 2017
- Oracle Database 12c Release 1
- Oracle Database 19c
- Oracle MySQL 8.0
- PostgreSQL 9.6.15 and 11.5

OAuth client configuration and persistent grants

- PingDirectory 7.0, 7.2, 7.3, 8.0, 8.1
- Microsoft Active Directory 2016
- Oracle Directory Server Enterprise Edition 11g
- Oracle Unified Directory 12c
- Amazon Aurora (MySQL 5.6.10a)
- Amazon Aurora (PostgreSQL 10.7)
- Microsoft SQL Server 2016 SP2 and 2017
- Oracle Database 12c Release 1
- Oracle Database 19c
- Oracle MySQL 8.0
- PostgreSQL 9.6.15 and 11.5
- Custom implementation through the PingFederate SDK

Registration and profile management of local identities

PingDirectory 7.0, 7.2, 7.3, 8.0, 8.1

Persistent authentication sessions

- PingDirectory 7.2, 7.3, 8.0, 8.1
- Amazon Aurora (MySQL 5.6.10a)
- Amazon Aurora (PostgreSQL 10.7)
- Microsoft SQL Server 2016 SP2 and 2017
- Oracle Database 12c Release 1 and 19c
- Oracle MySQL 8.0
- PostgreSQL 9.6.15 and 11.5



Note:

PingFederate was tested with vendor-specific JDBC drivers. For more information, see Database driver information on page 115.

Third-party cryptographic solutions

Hardware security modules (optional)

AWS CloudHSM

Client software version: 3.1.1



PingFederate must be deployed on one of the Linux operating systems supported by both AWS CloudHSM and PingFederate.

Gemalto SafeNet Luna Network HSM 6

HSM firmware version: 6.3

Firmware version: 6.27.0 Client software version: 6.3

Gemalto SafeNet Luna Network HSM 7

Appliance software version: 7.2.0

Firmware version: 7.2.0

Client software version: 7.2.0

nCipher nShield Connect (in FIPS 140-2 Level 3 mode)

Host and Firmware version: 12.40.0

Client driver version: 12.40.2

Hardware Model: Net HSM 6000 appliance



Note:

When integrating with a hardware security module (HSM), you must deploy PingFederate with Oracle Server JRE (Java SE Runtime Environment) 8 or Amazon Corretto 8.

Software cryptographic solution

Bouncy Castle

Bouncy Castle Java FIPS 1.0.2

Hardware requirements

Minimum hardware recommendations

- Multi-core Intel Xeon processor or higher
 - 4 CPU/Cores recommended
- 4 GB of RAM
 - 1.5 GB available to PingFederate
- 1 GB of available hard drive space



Note:

Although it is possible to run PingFederate on less powerful hardware, the guidelines provided accommodate disk space for default logging, auditing profiles, and CPU resources for a moderate level of concurrent request processing.

Database driver information

PingFederate is compatible with the following vendor-specific JDBC drivers.

PingFederate and JDBC driver compatibility

Database server	Driver information
Amazon Aurora (MySQL	Driver version information
5.6.10a)	mysql-connector-java version 8.0.19
	Driver class
	com.mysql.cj.jdbc.Driver
	JDBC URL
	<pre>jdbc:mysql://databaseservername/databasename</pre>
	Database location
	Regional
	Database features
	One writer and multiple readers
Amazon Aurora	Driver version information
(PostgreSQL 10.7)	postgresql version 42.2.5
	Driver class
	org.postgresql.Driver
	JDBC URL
	<pre>jdbc:postgresql://databaseservername/databasename</pre>
	Database features
	One writer and multiple readers
Microsoft SQL Server	Driver version information
2016 SP2 and 2017	sqljdbc version 7.2.1
	Driver class
	com.microsoft.sqlserver.jdbc.SQLServerDriver
	JDBC URL
	<pre>jdbc:sqlserver://databaseservername;databaseName=databa</pre>

Database server	Driver information		
Oracle Database 12c Release 1 and 19c	Driver version information		
Release Fallu 190	ojdbc7 version 12.1.0.2.0		
	Driver class		
	oracle.jdbc.OracleDriver		
	JDBC URL		
	jdbc:oracle:thin:@databaseservername:databasename		
Oracle MySQL 8.0	Driver version information		
	mysql-connector-java version 8.0.15		
	Driver class		
	com.mysql.cj.jdbc.Driver		
	JDBC URL		
	<pre>jdbc:mysql://databaseservername/databasename</pre>		
PostgreSQL 9.6.1 and	Driver version information		
11.2	postgresql version 42.2.5		
	Driver class		
	org.postgresql.Driver		
	JDBC URL		
	<pre>jdbc:postgresql://databaseservername/databasename</pre>		

For additional information about these drivers, contact the respective vendors.

Port requirements

The following table summarizes the ports and protocols that PingFederate uses to communicate with external components. This information provides guidance for firewall administrators to ensure the correct ports are available across network segments.



Direction refers to the direction of the initial requests relative to PingFederate. Inbound refers to requests PingFederate receives from external components. Outbound refers to requests PingFederate sends to external components.

PingFederate required ports and protocols

Service	Protocol, direction, transport, default port	Source	Destination	Description
Administrative console	HTTPS, inbound, TCP, 9999	Browsers accessing the administrative console, REST calls to the administrative API, web service calls to the Connection Management Service.	Administrative node	Used for incoming requests to the administrative console. Configurable in the run.properties file.
		Applicable to the console node in a clustered PingFederate environment.		
Administrative console	HTTPS, outbound, TCP, 443	Administrator accessing online help. Applicable to the console node in a clustered PingFederate environment.	docs.pingidentity	dosed for accessing online help from the administrative console.
Runtime engine	HTTPS, inbound, TCP, 9031 (and 9032 if configured)	Browsers accessing the runtime server for SSO or SLO; web service calls to the SSO Directory Service; REST calls to the OAuth Client Management Service, the OAuth Access Grant Management Service, the Persistent Grant Management API, and the Session Revocation API. Applicable to all runtime engine nodes in a clustered PingFederate environment.	Runtime engine nodes	Used for incoming requests to the runtime engine. Configurable in the run.properties file.

Service	Protocol, direction, transport, default port	Source	Destination	Description
Cluster traffic	JGroups, inbound, TCP, 7600	PingFederate peer servers in a clustered PingFederate environment.	Administrative node and runtime engine nodes	Used for communications between engine nodes in a cluster when the transport mode for cluster traffic is set to TCP (the default behavior). Configurable in the run.properties file.
Cluster traffic	JGroups, inbound, TCP, 7700	PingFederate peer servers in a clustered PingFederate environment.	Administrative node and runtime engine nodes	Used by other nodes in the cluster as part of the cluster's failure-detection mechanism when the transport mode for cluster traffic is set to TCP (the default behavior). Configurable in the run.properties file.
Cluster traffic (if configured)	JGroups, outbound, TCP, 443	PingFederate peer servers in a clustered PingFederate environment.	Amazon Simple Storage Service (Amazon S3) or an OpenStack Swift server	Used by all nodes when the optional dynamic discovery mechanism is enabled.
Cluster traffic	JGroups, inbound,UDP, 7601	PingFederate peer servers in a clustered PingFederate environment.	Administrative node and runtime engine nodes	Used for communications between engine nodes in a cluster when the transport mode for cluster traffic is set to UDP. By default, the transport mode is TCP. Configurable in the
				run.properties file.
PingOne for Enterprise integration (if configured)	HTTPS and secure WebSocket, TCP, 443	PingFederate Applicable to the console node in a clustered PingFederate environment.	pingone.com	Used for communications between PingFederate and PingOne for the purpose of establishing and maintaining a managed SP connection to PingOne for Enterprise, monitoring of PingFederate from the PingOne admin portal, authenticating end users against the PingOne for Enterprise Directory.
Active Directory domains/ Kerberos realms (if configured)	Kerberos, outbound, TCP or UDP, 88	PingFederate	Windows domain controllers	Used for communications between PingFederate and Windows domain controllers for the purpose of Kerberos authentication.

Service	Protocol, direction, transport, default port	Source	Destination	Description
reCAPTCHA (if configured)	HTTPS, outbound, TCP, 443	PingFederate	www.google.com recaptcha/api/ site verify	/Used by the HTML Form Adapter when invisible reCAPTCHA from Google is enabled to prevent automated attacks.
Administration notification	SMTP, outbound, TCP, 25 (465 if SMTPS)	All nodes	SMTP server	Used to send notification messages for various events. For more information, see <i>Runtime notifications</i> on page 863.



For PingID integration, see PingID required domains, URLs, and ports

Depending on the integration kits deployed and the connecting third-party systems, such as email server or SMS service provider, additional ports may be required.

Installing Java

PingFederate requires a Java runtime to be installed on your server.

About this task

PingFederate has been tested in the following Java environments:

- Amazon Corretto 11
- Amazon Corretto 8
- OpenJDK 11
- Oracle Java SE Development Kit 11 LTS
- Oracle Java SE Runtime Environment (Server JRE) 8



Ping Identity Java Support Policy applies. For more information, see Java Support Policy in the Ping Identity Knowledge Base.



Important:

Due to the import restrictions of some countries, Oracle Server Java SE Runtime Environment (JRE) 8 has built-in restrictions on available cryptographic strength (key size). To use larger key sizes, enable the Java Cryptography Extension (JCE) unlimited strength jurisdiction policy. For more information, see the Java 8 release notes in Oracle's documentation.

For Oracle Java SE Development Kit 11, the JCE jurisdiction policy defaults to unlimited strength. For more information, see the Oracle JDK Migration Guide in Oracle's documentation.

Steps

- 1. Download and install a Java runtime.
- 2. Set the JAVA_HOME environment variable to the Java installation directory path and add its bin directory to the PATH environment variable.



Note:

If you intend to use the PingFederate installer for Windows or run PingFederate as a service, you must set the JAVA HOME environment variable and modify the PATH environment variable at the system level. If you are not using the PingFederate installer or running PingFederate as a service, you can set the variables at either the system or user level.



CAUTION:

When running PingFederate for Windows, switching the Java version from 8 to 11 (or the reverse) will prevent the service from running, and you will not be able to start PingFederate. The problem occurs because garbage collection logging configuration arguments that are used by Java 8 are incompatible with those used by Java 11.

To change Java versions:

- a. Run <pf install>\pingfederate\sbin\win-x86-64\uninstall-service.bat to deregister the PingFederate service.
- b. Install the new Java version and update the JAVA HOME and PATH environment variables.
- c. Run <pf install>\pingfederate\sbin\win-x86-64\install-service.bat to register the PingFederate service.

Installation options

PingFederate installs on Windows and Linux operating systems.

Install PingFederate using the following methods:

- Install PingFederate on a Windows system by running the installer for Windows or by extracting the distribution . zip file. Using the installer for Windows is the preferred method.
- Install PingFederate on a Linux system by extracting the distribution .zip file.



This documentation refers to the installation directory path as <pf_install>. This is where the pingfederate directory is located. For example, opf install/pingfederate/bin.



Important:

To avoid future problems with automated upgrades, do not rename the installed pingfederate directory.

If you are installing multiple instances of PingFederate on the same machine, for example, a console node and an engine node in a clustered environment, install each instance using a unique <pf install> directory.

If you are upgrading an existing PingFederate environment, see Upgrading PingFederate in the PingFederate Server documentation.

Installing PingFederate on Windows

You can install PingFederate on a Windows system using the installer for Windows or the distribution .zip file. Using the installer for Windows is the preferred method.

Before you begin

- Request a license key through the Ping Identity licensing page.
- Ensure your are logged on to your system with sufficient privileges to install and run an application.
- Verify that you have installed the Java runtime and that you have set the required environment variables correctly. For more information, see *Installing Java* on page 119 in the PingFederate Server documentation.

Steps

1. Install PingFederate using the installer for Windows or the distribution .zip file as described in the following table.

Installation medium	Steps		
PingFederate installer for	To install PingFederate using the installer for Windows:		
Windows	 Download the PingFederate installer for Windows from the Ping Identity website. 		
	b. Double-click the .msi file to open the PingFederate Setup Wizard, and follow the instructions to complete the installation.		
	PingFederate is configured to run as a service and starts automatically at the end of the installation process.		
	Note:		
	The PingFederate installer for Windows installs only one instance of PingFederate on a Windows server. If you need additional PingFederate instances on the same Windows server, install them using the distribution .zip file.		
	You must manually configure various port settings in the <pf_install>/pingfederate/bin/run.properties file each instance to avoid port conflicts.</pf_install>		
Distribution .zip file	To install PingFederate using the distribution .zip file:		
	 a. Download the distribution .zip file from the Ping Identity e. The distribution .zip file is identical for both Windows and Linux. b. Extract the file into an installation directory. 		

2. If you have installed PingFederate by extracting the distribution .zip file, start PingFederate manually by running the following script:

```
<pf install>/pingfederate/bin/run.bat
```

Wait for the script to finish—the startup process completes when this message appears near the end of the sequence:

PingFederate running...



To configure PingFederate to run as a service, follow the steps in *Installing PingFederate service on* Windows manually on page 126 in the PingFederate Server documentation.



CAUTION:

When running PingFederate for Windows, switching the Java version from 8 to 11 (or the reverse) will prevent the service from running, and you will not be able to start PingFederate. The problem occurs because garbage collection logging configuration arguments that are used by Java 8 are incompatible with those used by Java 11.

To change Java versions:

- a. Run <pf install>\pingfederate\sbin\win-x86-64\uninstall-service.bat to deregister the PingFederate service.
- b. Install the new Java version and update the JAVA_HOME and PATH environment variables.
- c. Run <pf install>\pingfederate\sbin\win-x86-64\install-service.bat to register the PingFederate service.

Result

If your organization requires compliance with FIPS 140-2 or plans on managing keys and certificates using a hardware security module (HSM), see Supported hardware security modules on page 169 in the Getting Started section of the PingFederate Server documentation.

Installing PingFederate on Linux systems

You install PingFederate on a Linux system using the distribution .zip file.

Before you begin

- See System requirements for a list of qualified Linux operating systems.
- Request a license key through the Ping Identity licensing website.
- Ensure you are logged on to your system with sufficient privileges to install and run an application. You must install and run PingFederate under a local user account.
- Verify that you have installed the Java runtime and that you have set the required environment variables correctly. For more information, see *Installing Java* on page 119.

Steps

- 1. Download the latest version of the PingFederate Server distribution .zip file from the Ping Identity website.
- 2. Extract the file into the target installation directory.

3. Start PingFederate manually by running the following script.

```
<pf install>/pingfederate/bin/run.sh
```

The startup process completes when this message appears near the end of the sequence.

```
PingFederate running...
```



Tip:

To configure PingFederate to run as a service, follow the steps in *Installing PingFederate service on* Linux manually.

Result

If your organization requires compliance with FIPS 140-2 or plans on managing keys and certificates using a hardware security module (HSM), see Supported hardware security modules.

Installing the PingFederate service on Linux manually

If you have not installed PingFederate on Linux using the distribution .zip file, you can install it manually.

Before you begin

- Request a license key through the Ping Identity licensing website.
- Ensure you are logged on to your system with sufficient privileges to install and run an application.
- Verify that you have installed the Java runtime and that you have set the required environment variables correctly. For more information, see *Installing Java* on page 119 in the PingFederate Server documentation.

About this task

To install the PingFederate service on Linux manually:

Steps

- 1. Download the distribution .zip file from the Ping Identity website.
- 2. Extract the file into an installation directory, <pf install>.
- 3. Create a new local user account for the PingFederate service; for example, pingfederate.



Note:

The service account is referred to as <pf user>.

4. Change the ownership of the PingFederate installation directory <pf install> and update the read and write permissions using the following commands.

```
chown -R <pf user> <pf install>
chmod -R 775 < pf install>
```

- 5. If the operating system supports systemd, follow these steps to install the PingFederate unit file.
 - a. Edit the pingfederate.service systemd unit file, located in the <pf install>/ pingfederate/sbin/linux directory.

Replace the following variables with information from your environment:

```
${PF VERSION}
```

The version of PingFederate.

```
${PF_USER}
```

The local user account for the PingFederate service.

```
${PF HOME}
```

The <pf install>/pingfederate directory.

For example, if <pf install> is /opt/identity.fed, replace \${PF_HOME} with / opt/identity.fed/pingfederate.

```
${PF_JAVA_HOME}
```

The JAVA HOME environment variable value (a directory).

b. Copy the pingfederate.service file to the systemd unit files directory. For example, /etc/ systemd/system.



Note:

Depending on the operating system, the exact location might vary. Consult your system administrators as needed. The rest of the step assumes /etc/systemd/system is the systemd unit files directory.

c. Use the following command to update the read and write permissions of the pingfederate.service systemd unit file.

```
chmod 664 /etc/systemd/system/pingfederate.service
```

d. Use the following commands to load the new system configuration changes and start the PingFederate service.

```
systemctl daemon-reload ; \
systemctl start pingfederate
```

e. Use the following commands to configure the PingFederate service to start automatically as the server boots.

```
systemctl enable pingfederate ;\
systemctl daemon-reload ;\
systemctl restart pingfederate
```

Result:

After setting up the PingFederate systemd unit file, you can use the following systemctl command to manage the PingFederate service.

```
systemctl start pingfederate
systemctl stop pingfederate
systemctl restart pingfederate
systemctl status pingfederate
```

- 6. If the operating system supports SysV initialization, follow these steps to install the PingFederate script.
 - a. Edit the pingfederate script, located in the <pf install>/pingfederate/sbin/linux directory.

Replace the following statements with information from your environment:

```
PF HOME=$PF HOME
```

Replace \$PF_HOME with the <pf install>/pingfederate directory.

For example, if cpf_install>is /opt/identity.fed, replace \$PF_HOME with /opt/identity.fed/pingfederate.

USER="pingfederate"

If the PingFederate service account is not pingfederate, replace pingfederate with the local user account for the PingFederate service.

For example, if <pf user> is pingfed, replace pingfederate with pingfed.

Example:

Example (truncated)

If <pf_install> and <pf_user> are /opt/identity.fed and pingfederate
respectively, the required modifications are as follows

```
PF_HOME=/opt/identity.fed/pingfederate
DIR="$PF_HOME/sbin"
USER="pingfederate"
...
```

Copy the pingfederate script to the SysV initialization directory; for example, /etc/rc.d/init.d.

The exact location might vary, depending on the operating system. Consult your system administrators, as needed. The rest of the step assumes /etc/rc.d/init.d is the SysV initialization directory.

c. Use the following command to update the read and write permissions of the pingfederate SysV initialization script.

```
chmod 755 /etc/rc.d/init.d/pingfederate
```

d. Configure the operating system to start the PingFederate service at various runlevels.

On an RHEL server, you can use the Service Configuration utility to do so.

Alternatively, the initialization directories associated with various runlevels can accept manual symbolic links of the pingfederate script using the ln -s source target command.

Example:

You can create the following symbolic links on an RHEL server where runlevels 2 and 4 are not used.

```
ln -s /etc/rc.d/init.d/pingfederate /etc/rc3.d/S84pingfederate
ln -s /etc/rc.d/init.d/pingfederate /etc/rc5.d/S84pingfederate
ln -s /etc/rc.d/init.d/pingfederate /etc/rc0.d/K15pingfederate
ln -s /etc/rc.d/init.d/pingfederate /etc/rc1.d/K15pingfederate
ln -s /etc/rc.d/init.d/pingfederate /etc/rc6.d/K15pingfederate
```

Some operating systems might require a restart of the system to activate the new scripts. Consult your system administrators as needed.

Result:

After setting up the PingFederate SysV initialization script, you can use the **Service**Configuration utility or the following **service** commands to manage the PingFederate service.

```
service pingfederate start
service pingfederate stop
service pingfederate restart
service pingfederate status
```

Installing PingFederate service on Windows manually

If you have not installed PingFederate using the installer for Windows, you can install the PingFederate service manually.

Before you begin

- Request a license key through the Ping Identity licensing website
- Ensure you are logged on to your system with sufficient privileges to install and run an application.
- Verify that you have installed the Java runtime and that you have set the required environment variables correctly. See *Installing Java* on page 119 in the PingFederate Server documentation.

About this task

If you have installed PingFederate using the installer for Windows, skip these steps, because PingFederate has already been configured to run as a service and to start automatically at the end of the installation process.

Steps

- 1. Download the distribution .zip file from the Ping Identity *website*. The distribution .zip file is identical for both Windows and Linux.
- 2. Extract the file into an installation directory, <pf install>.
- 3. Start PowerShell or Command Prompt as an administrator.
- 4. Run the <pf install>\pingfederate\sbin\win-x86-64\install-service.bat file.
- 5. Open the Control Panel # Administrative Tools # Services management console.
- 6. Right-click the **PingFederate** service and select **Start**.

Result

PingFederate service starts automatically on reboot.

Uninstalling PingFederate

This section explains how to uninstall PingFederate from a Windows or Linux server.

Uninstalling PingFederate involves removing the previously-installed PingFederate service and the installation directory, <pf install>.

Uninstalling PingFederate from a Windows server

The method you use to uninstall PingFederate depends on whether you installed it using the installer for Windows or the distribution .zip file.

Before you begin

- Ensure you are logged on to your system with sufficient privileges to uninstall an application.
- Optionally, make a backup copy of the PingFederate installation directory <pf install>.

Steps

1. Verify the installation medium in **Control Panel** # **Uninstall a Program**.



Note:

The existence of a PingFederate entry indicates a previous installation of PingFederate using the Windows installer. If there is no existing PingFederate entry, the distribution .zip file was used to perform the installation.

- 2. Depending on the type of installation, uninstall PingFederate using one of the following methods. Choose from:
 - PingFederate installer for Windows

Use Control Panel # Uninstall a Program to uninstall PingFederate. Doing this removes the PingFederate service and the installation directory.

- Distribution .zip file
 - a. Open the Control Panel # Administrative Tools # Services management console.
 - b. Right-click the **PingFederate** service and select **Stop**.
 - c. Run uninstall-service.bat from the <pf install>\pingfederate\sbin subdirectory that corresponds to your platform processor.
 - d. Optionally, remove the PingFederate installation directory <pf install>.

Uninstalling PingFederate from a Linux server

You can use the systemd service or the SysV initialization script to uninstall PingFederate from a Linux server.

Before you begin

Ensure you are logged on to your system with sufficient privileges to uninstall an application.

Steps

- 1. Uninstall PingFederate using one of the following methods. Choose from:
 - PingFederate systemd service

Use the following systemct1 commands to stop and disable Pingfederate.

```
systemctl stop pingfederate ;\
systemctl disable pingfederate ;\
```

systemctl daemon-reload

You can also remove the PingFederate systemd unit file pingfederate.service from the systemd unit files directory/etc/systemd/system prior to running the systemctl daemonreload command.

PingFederate SysV initialization script

Do one of the following:

- Use the Service Configuration utility to stop and disable the PingFederate service.
- Remove any symbolic links from various initialization directories to stop the PingFederate service.
- Remove the PingFederate SysV initialization script pingfederate from the SysV initialization directory /etc/rc.d/init.d.



Note:

Depending on the operating system, the exact directory locations may vary. Consult your system administrators as needed.

2. Optional: Remove the PingFederate installation directory, <pf install>.

Upgrading PingFederate

The PingFederate Upgrade Guide describes how to upgrade your PingFederate environment to the latest version on Windows and Linux systems.

Depending on your PingFederate installation, you can upgrade by using the PingFederate installer for Windows or the Upgrade Utility, which automatically migrates existing PingFederate installations of version 6.0 and later to version 10.1. As of PingFederate 10.0, the Upgrade Utility is included with the software distribution rather than being a separate download.

For general information about upgrading Ping Identity products and an upgrade planning guide, see Planning your upgrade.



Important:

We recommend that you upgrade your test environment and verify that the new installation meets your expectations before upgrading your production environment. Also thoroughly retest the behavior of any customized components. After you complete the upgrade process, you can create a backup of your previous installation and remove it from the server.

Additionally, end users might experience service disruptions as you upgrade your PingFederate environment. As needed, schedule a maintenance window to perform the upgrade.

Product downloads

Download the latest version of PingFederate from the *PingFederate Downloads* website.

- To download the PingFederate installer for Windows, click Windows.
- To download the PingFederate product distribution .zip file, click Linux. The distribution .zip file can be used to upgrade PingFederate on both Windows and Linux.

Upgrade paths

Operating system	Source version	Source installation medium	Possible upgrade paths
Microsoft Windows	8.x through 10.x	PingFederate installer for Windows	PingFederate installer for Windows, or PingFederate Upgrade Utility
	6.x through 10.x	PingFederate product distribution . zip file	PingFederate Upgrade Utility
Linux	6.x through 10.x	PingFederate product distribution . zip file	PingFederate Upgrade Utility

If you are upgrading from PingFederate 5.3 or earlier, contact Support for more information.

Both the PingFederate installer for Windows and the Upgrade Utility create a new installation based on the new product distribution . zip file, and then copy the relevant files and property values from the existing installation (the source) to the new installation (the target). As a result, neither tool affects the source installation.

Integration kits

Both upgrade tools also copy the program files for the deployed adapters, connectors, and token translators (the integration kits in general) from the source installation to the target installation. While the tools do not upgrade the integration kits automatically, you can download newer versions by visiting the Ping Identity Downloads website and upgrade the integration kits manually. Documentation for integration kits is available on the Ping Identity website.

Preparation

Prepare for the PingFederate upgrade by:

- Reviewing the PingFederate release notes for enhancements, upgrade considerations, deprecated features, and other known issues and limitations.
- Reviewing the post-upgrade tasks.
- Reviewing potential changes in system and port requirements.
- Obtaining a new license key if needed.
- Updating the Java runtime to version 8 or 11 on your PingFederate servers if needed.

When you upgrade the Java runtime, modify the previously defined paths for the system JAVA_HOME and PATH environment variables.



Important:

PingFederate versions prior to 7.2 will not start using the currently supported Java runtime. If you need to start the previous PingFederate version on the same server after the upgrade, retain the older Java installation and change environment variables back when needed.

 Completing any unfinished connections (Drafts) in the administrative console, if you want to include them in the migration.

Upgrade considerations

Several specific modifications since version 10.0 might affect existing deployments.

Expression Admin role

When upgrading to PingFederate 10.1 from a previous version, administrative users who were granted the Admin role in the earlier installation are granted the Expression Admin role automatically. You can achieve the same result by using the /bulk/import administrative API endpoint to bulk-import a configuration that was bulk-exported from PingFederate 10.0.

Additionally, all four administrative roles, namely User Admin, Admin, Expression Admin, and Crypto Admin, are required to access and make changes through the following services:

- The /bulk, /configArchive, and /configStore administrative API endpoints
- The System # Server # Configuration Archive window in the administrative console
- The Connection Management configuration item on the Security # System Integration # Service Authentication window

Authentication session created after user registration

As of PingFederate 10.1, an authentication session is automatically created for a user after registration, preventing the user from having to log in again during the next SSO transaction. This feature is enabled by default for all new and existing local identity profiles. However, if needed, you can disable it through the /localIdentity/identityProfiles administrative API endpoint by setting the createAuthnSessionAfterRegistration attribute to false.

AWS CloudHSM

If PingFederate is running on Linux and uses AWS CloudHSM, when administrators upgrade from PingFederate version 10.0 or earlier to PingFederate version 10.0.1 or later, they must also upgrade the CloudHSM client to version 3.1.1.

Template html.form.login.template.html

Starting with PingFederate 10.0, the html.form.login.template.html template no longer includes the \$forgotPasswordUrl variable.

Upgrade considerations introduced in PingFederate 9.x

Gemalto SafeNet Luna HSM 6.3

When integrating with Gemalto SafeNet Luna Network HSM 6 (hardware security module), PingFederate 9.2 requires firmware version of 6.3.0 and client driver version of 6.3. See *Integrating with Gemalto SafeNet Luna Network HSM* on page 173 for setup information.

Weaker cipher suites disabled

Starting with PingFederate 9.1, weaker cipher suites TLS_RSA_WITH_AES_128_CBC_SHA and TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA are disabled in new installations and upgrades. As a result, the administrative and runtime servers support only TLS 1.2. If you must re-enable these cipher suites for legacy clients, refer to *Managing cipher suites* on page 820 for more information.

LDAP service accounts on PingDirectory

If PingFederate 9.3.1 or newer has an LDAP connection with PingDirectory, then add the configread privilege to its service account in PingDirectory. Otherwise, users will not receive password expiry notifications. For more information, see *Assigning Privileges to Normal Users and Individual Root Users* in the PingDirectory documentation.

Improved validation for AudienceRestriction

If an IdP connection is configured with multiple virtual server IDs, the AudienceRestriction value in a SAML response must now match the virtual server ID information embedded in the protocol endpoint at which PingFederate receives the message. Otherwise the SSO attempt fails. To override this validation on a per-connection basis, see Configuring validation for the AudienceRestriction element on page 834.

Custom authentication selector

If you have created a custom authentication selector that returns an IdP adapter instance ID or the connection ID of an IdP connection, you must update the associated descriptor instance. See *Updating the custom authentication selector* on page 152 for more information.

Provisioning datastore reset

Upgrading to PingFederate 9.0 or 9.0.1 when using its outbound provisioning capability can result in user records being disabled at SaaS applications. The issue is resolved in version 9.0.2.

If you are upgrading from version 8.4.4 (or earlier) or from version 9.0.2, 9.0.3, and 9.0.4 to version 10.0, the upgrade process automatically resolves this issue. No further action is required.

If you are upgrading from version 9.0 or 9.0.1 to PingFederate 10.0, you must use the proving command-line tool to reset the provisioning datastore on the upgraded installation. See *Reviewing database changes* on page 146 for more information.

Security enhancement in JDBC datastore queries

A security enhancement has been made in PingFederate 9.0 to safeguard JDBC datastore queries against back-end SQL injection attacks. This protection is enabled for all new installations. For upgrades, see *Reviewing database changes* on page 146.

Access token validation response

Starting with PingFederate 9.2, the access token validation response no longer includes the username and subject elements by default. Responses include them only if they were mapped in the issuing access token management instance.

Upgrade considerations introduced in PingFederate 8.x

InterReqStateMgmtMapImpl.expiry.mins renamed

The InterReqStateMgmtMapImpl.expiry.mins setting in the size-limits.conf file has been renamed in PingFederate 8.4.2. If you have previously modified the value of this setting, please refer to Copying customized files or settings on page 144 for more information.

An improved index (IDX FIELD NAME) in the database table for OAuth clients

PingFederate 8.4 has modified an existing index (IDX_FIELD_NAME) in the pingfederate_oauth_clients_ext database table as a general improvement. For information about modifying this index in your existing table, see *Reviewing database changes* on page 146.

Security enhancement to the OAuth token endpoint

Starting with version 8.3, a new PingFederate installation no longer allows OAuth clients to send access token validation requests to its token endpoint (/as/token.oauth2) by the HTTP GET method.

For upgrades, the Upgrade Utility applies this new behavior to the new installation as well unless the $<\!pf_install>$ /pingfederate/server/default/data/config-store/oauth-token-endpoint-binding.xml file has been modified in the older version, in which case the Upgrade Utility preserves the modified configuration.

SSLv2Hello disabled

Starting with PingFederate 8.3, SSLv2Hello is disabled. Customers are encouraged to update their applications to use TLSv1, TLSv1.1, or TLSv1.2 when establishing HTTPS connections with PingFederate.

As needed, SSLv2Hello can be re-enabled as needed. See *Enabling SSLv2Hello* for more information.

License management simplification

Starting with version 8.2, PingFederate no longer maintains its license information in the $< pf_install > / pingfederate / server / default / data / .pingfederate . lic file, which is known as the secondary license file in the previous versions of PingFederate. The .pingfederate.lic, if any, is ignored.$

We recommend using the administrative console to simplify the license management aspect of a standalone PingFederate server or a clustered PingFederate environment.

Security enhancement for a clustered PingFederate environment

As of PingFederate 8.1, when encryption is enabled for the network traffic sent between nodes in a clustered PingFederate environment, you must provide an authentication password for the cluster as well; otherwise PingFederate aborts during its startup process.

For more information about the pf.cluster.encrypt and pf.cluster.auth.pwd properties, see *Deploying cluster servers* on page 201.

Metadata signing

Previously, when no signing certificate was chosen on the Metadata Signing tab on the System # Protocol Metadata # Metadata Settings window, the /pf/sts_mex.ping and /pf/federation_metadata.ping system-services endpoints provided signed WS-Trust and WS-Federation metadata using one of the certificates configured on the Security # Certificate & Key Management # Signing & Decryption Keys & Certificates window.

Starting with PingFederate 8.1, if no certificate is selected in the **Metadata Signing** menu, PingFederate provides unsigned metadata at both aforementioned endpoints. Select a certificate in the **Metadata Signing** window if signed metadata is desired.

Hostname verification for email server

For email notification using SSL or TLS, hostname verification of the certificate is available starting with PingFederate 8.1. This option is enabled automatically when the **Use SSL** or **Use TLS** check box is selected for a new configuration. When upgrading from a previous version of PingFederate, if email notification had already been configured to use SSL or TLS, the Upgrade Utility preserves the configuration without activating the hostname verification option for compatibility reasons. Administrators should consider activating this new option to improve security.

New login template file for the HTML Form Adapter

Previously, when multiple instances of the HTML Form Adapter are chained together (for example, in an instance of the Composite Adapter), the subsequent instance tried authenticating the end user with the credentials from the previous login, which might fail when the HTML Form Adapter instances were configured to use different password credential validators (PCVs). Although this use case is rare, PingFederate 8.1 has corrected the behavior. As a result, the login template file, $<\!pf_install\!>$ /pingfederate/server/default/conf/template/html.form.login.template.html, has been modified.

If you have previously customized this login template file and if you have authentication use cases that chain multiple instances of the HTML Form Adapter, you should re-customize using the new html.form.login.template.html file.

New connection pool library

As of PingFederate 8.0, support for BoneCP as the JDBC connection pool library has been deprecated and replaced with Apache Commons DBCP 2, which requires JDBC 4.1 or later drivers.

Verify the database-driver JAR files, found in the $<\!pf_install>$ /pingfederate/server/default/lib directory, meet the minimum version requirement. If you are using JDBC drivers of version 4.0 (or earlier), contact your vendors for the latest drivers and replace the older JDBC database-driver JAR files with the latest.

For more information, including re-enabling BoneCP as the JDBC connection pool library, see *Reviewing database changes* on page 146.

Log4j 2 upgrade

PingFederate 8.0 has upgraded its logging framework from Log4j to Log4j 2.

If you have previously customized <pf install>/pingfederate/server/default/conf/ log4j.xml, you will need to manually migrate your changes to the new log4j2.xml in the same conf directory. See Reviewing log configuration on page 151 for instructions.



Note:

PingFederate was tested with vendor-specific JDBC drivers. For more information, see *Database* driver information on page 115. To obtain the database driver JAR file, contact your database vendor. Install the database driver file to the <pf install>/pingfederate/server/default/ lib directory then restart the server.

Upgrade considerations introduced in PingFederate 7.x

Hostname verification for LDAPS

For LDAP type datastores with LDAPS enabled, hostname verification of the certificate is enabled by default for all new datastores beginning with PingFederate 7.3. When upgrading from a previous version of PingFederate, this option is disabled for existing datastores for compatibility reasons. Administrators should consider activating this new option for greater security.

Changes in a database table supporting nested group membership

Outbound provisioning of groups and nested group membership requires an update in the internal datastore. Follow the instructions in Reviewing database changes on page 146 to add or update the group_membership table.

SSLv3 disabled

To mitigate the POODLE attack, the SSLv3 protocol is disabled by default starting in PingFederate 7.3. It can be re-enabled by modifying the connector configuration in jetty-runtime.xml and jetty-admin.xml found in the <pf install>/pingfederate/etc directory.

New representation for multivalued attributes in WS-Federation assertions

Starting with PingFederate 7.3, multivalued attributes in WS-Federation assertions are now represented as multiple AttributeValue elements under a single Attribute element. Previously, they were represented as a series of Attribute elements with the same name. The new behavior was implemented for compatibility with ADFS 2.0. To revert to the previous behavior, a setting is available in wstrust-global-settings.xml.

A new index (EXPIRESIDX) in the database table for OAuth persistent grants

PingFederate 7.3 added an index (EXPIRESIDX) for the expires column in the pingfederate access grant database table. For information on adding this index to your existing table, see Reviewing database changes on page 146.

A new database table for OAuth persistent grant extended attributes

Starting with PingFederate 7.2 R2, a new database table needs to be created to support OAuth's persistent grant extended attributes. The database scripts to create this table can be found in <pf install>/pingfederate/server/default/conf/access-grant/sql-scripts/ access-grant-attribute-<databaseServer>.sql. See Reviewing database changes on page 146for more information.

LDAP filter syntax checking

Starting with PingFederate 7.2, LDAP filters only allow spaces in matched-against values.

Examples

```
(|(sAMAccountName=${username})(employeeID=ID for ${username})) is allowed;
spaces in the matched-against value of "ID for ${username}" are valid.
```

(| (sAMAccountName=\${username}) (employeeID=ID for \${username})) is not allowed because this filter contain spaces outside of matched-against values.

Invalid filters cause SSO runtime failures. Error messages logged to server.log include:

```
Caused by: javax.naming.NamingException: [LDAP: error code 87 -
Expected a closing parenthesis...
Caused by: javax.naming.NamingException: [LDAP: error code 87 -
 Unexpected closing parenthesis found...
```

We recommend reviewing LDAP filters and removing spaces outside of matched-against values after upgrade.

HTML Form Adapter enhancement

Starting with version 7.1 R3, PingFederate tracks login attempts in the HTML Form Adapter. When the number of login failures reaches the Challenge Retries threshold defined in the adapter, the user is locked out for one minute. See HTML Form Adapter on page 287 for more information.

A new index (CLIENTIDIDX) in the database table for OAuth persistent grants

PingFederate 7.1 R3 added an index (CLIENTIDIDX) for the client id column in the pingfederate access grant database table. For information on adding this index to your existing table, see Reviewing database changes on page 146.

Requested (formerly SAML) AuthN Context authentication selector process order changed

In releases prior to 7.1 R2, when the Requested AuthN Context Authentication Selector received a list of authentication contexts, it used the last context that it could match, rather than the first. However, both the SAML and OpenID Connect specifications treat an authentication context list as appearing in order of preference. To align the Requested AuthN Context Authentication Selector with these specifications, the selection order was changed in 7.1 R2. With this release, the selector will use the first authentication context it can match, rather than the last.

Multivalued LDAP attributes passed to outbound provisioning OGNL expressions

In releases before version 7.1, if an OGNL expression was used to populate a SaaS-partner field in outbound provisioning, only the first value of a selected multivalued LDAP attribute was used in the OGNL expression. As of PingFederate 7.1, this behavior was changed to use all values in the expression.



Note:

If this new behavior conflicts with existing deployments, it can be reverted via the supportMultiValuesFromDirectory property located in the <pf install>/pingfederate/server/default/data/config-store/ com.pingidentity.provisioner.mapping.OgnlFieldMapper.xml file.

OAuth clients reconfiguration

Neither the Upgrade Utility nor the platform-specific installers migrates OAuth clients that are created from PingFederate 6.5 through 7.0. Use any of the following interfaces to reconfigure your OAuth clients:

- The Clients window (Applications # OAuth # Clients) in the PingFederate administrative console.
- The /oauth/clients administrative API endpoint.

 The REST-based web service for OAuth client management at the /pf-ws/rest/oauth/ clients and /pf-ws/rest/oauth/clients/id endpoints. This web service requires the client records to be stored in a database.

Note that PingFederate has been storing OAuth clients in XML files since version 7.1; these clients are migrated to the new installation. In addition, if you have configured PingFederate 6.8 (or a more recent version) to store OAuth clients in an external database, the new installation retains that configuration as well.

Upgrade considerations introduced in PingFederate 6.x

Cluster bind address required

Starting with PingFederate 6.11, the pf.cluster.bind.address property located in <pf install>/pingfederate/bin/run.properties is required when running PingFederate in a cluster. The default value is NON LOOPBACK.

Decryption and digital signing policy changes

Potential security vulnerabilities have resulted in the following changes to PingFederate as of version 6.11. In some cases, these might impact interoperability with partners:

 When acting as an SP and using the POST binding, PingFederate decrypts an assertion only when the SAML response has been signed. An unsigned SAML response that contains an encrypted assertion is rejected.



Note:

Although strongly discouraged, this policy change can be reverted on a per-connection basis via the EntityIdsToAllowAssertionDecryptionWithoutResponseSignature list located in the <pf install>/pingfederate/server/default/data/config-store/ org.sourceid.saml20.profiles.sp.HandleAuthnResponse.xml file.

When acting as an IdP, PingFederate always signs a SAML response (even when the assertion is also signed) if it contains an encrypted assertion.



Although strongly discouraged, this policy change can be reverted on a per-connection basis via the EntityIdsToOmitResponseSignatureOnSignedEncryptedAssertion list located in the <pf install>/pingfederate/server/default/data/config-store/ org.sourceid.saml20.profiles.idp.HandleAuthnReguest.xml file.

 When acting as an IdP, PingFederate decrypts an encrypted NameID in an Attribute Query only when the request has been signed or the client has authenticated with basic or mutual TLS.

Key transport algorithm deprecated

Due to security risks associated with the RSA-v1.5 algorithm used for key transport, it is no longer available for new connections starting with PingFederate 6.11. Existing connections in which this algorithm is configured continue to support it. However, we recommend upgrading to use the newer RSA-OAEP algorithm. See Selecting an encryption certificate on page 452 for SP connections and Choosing an encryption certificate (SAML 2.0) on page 737 for IdP connections.

OAuth persistent grants expiration

When upgrading to PingFederate 6.8 later, all persistent grants for any existing OAuth deployments using an Oracle MySQL database will expire. To address this issue, the expires column in the

pingfederate access grant table should be set to null prior to the upgrade. If necessary, contact Ping Identity support for assistance.

OAuth clients reconfiguration

Neither the Upgrade Utility nor the platform-specific installers migrates OAuth clients that are created from PingFederate 6.5 through 7.0. Use any of the following interfaces to reconfigure your OAuth clients:

- The Clients window (Applications # OAuth # Clients) in the PingFederate administrative console.
- The /oauth/clients administrative API endpoint.
- The REST-based web service for OAuth client management at the /pf-ws/rest/oauth/ clients and /pf-ws/rest/oauth/clients/id endpoints. This web service requires the client records to be stored in a database.

Note that PingFederate has been storing OAuth clients in XML files since version 7.1; these clients are migrated to the new installation. In addition, if you have configured PingFederate 6.8 (or a more recent version) to store OAuth clients in an external database, the new installation retains that configuration as well.

OGNL library upgraded

The OGNL library has been upgraded in version 6.4. If you use OGNL expressions in versions prior to 6.4, we recommend retesting the expressions using the PingFederate administrative console or runtime tests.

Updating to the latest maintenance release

For PingFederate maintenance releases, you have the option to update your installation using the incremental update package. This in-place update method lets you replace and merge only the files that have changed.

Before you begin

- Ensure your installation is running PingFederate 10.1.x, where "x" represents an older maintenance release of 10.1.
- Make a backup copy of the PingFederate home directory.

About this task

You can incrementally update your PingFederate 10.1 installation to the latest 10.1 maintenance release. For example, if you are on PingFederate 10.1.1 and the latest maintenance release is version 10.1.4, you can use the in-place update method to update your installation to version 10.1.4.

The in-place update doesn't contain files from Ping Identity integration kits. You can upgrade an integration kit manually by downloading the latest kit from the PingFederate Downloads page's Add-ons tab and following the instructions provided by the kit's documentation.



Important:

You cannot use the in-place update method to upgrade an older version of PingFederate, such as 9.3.3 or 10.0.2, to version 10.1.x. For those older versions, you must use the standard upgrade method for your platform described in *Upgrading PingFederate* on page 128.



Note:

The in-place update method involves manual modification of PingFederate files. If you're not comfortable moving and editing these files, then instead use the standard upgrade method for your platform described in *Upgrading PingFederate* on page 128.

If your installation includes a cluster, perform the following procedure on each node starting with the administrative console node.

Steps

- 1. Go to the *PingFederate Previous Releases* page on the Ping Identity website.
- 2. In the **PingFederate 10.1.x** section, download the maintenance in-place update and extract its contents from the .zip file.
- 3. Stop PingFederate.
- 4. Copy the files in the in-place update's pingfederate directory and paste them into their corresponding locations in your current PingFederate installation, replacing the old files. For example, if the in-place update contains pingfederate/server/default/lib/pf
 - protocolengine.jar, copy the pf-protocolengine.jar file to the same location in your PingFederate installation: install>/pingfederate/server/default/lib.
- 5. Compare each file in the in-place update's merge required directory with the version of the file in your current PingFederate installation and manually merge the changes into the file in your current installation.
 - For example, if the in-place update contains merge required/pingfederate/server/ default/conf/language-packs/pingfederate-messages.properties, copy the changes in the new version into the pingfederate-messages.properties file in your PingFederate installation.
- 6. Start PingFederate.

Upgrading PingFederate on Windows using the installer

If you used the PingFederate installer for Windows to install the previous version of PingFederate, you can use it to upgrade to the current version of PingFederate.

Before you begin

- Read the Upgrading PingFederate on page 128 topic for an overview of the upgrade process.
- Ensure that you are logged on to your server with appropriate privileges to install and run an application.

About this task

Upgrade results are contained in the upgrade.log file. If the upgrade succeeds, upgrade.log is located in <pf install target>\pingfederate\upgrade\log. If the upgrade fails, upgrade.log is located in <pf install target>. The default location for <pf install target> is C:\Program Files\Ping Identity. Administrators can change the location during the upgrade process.



Note:

The PingFederate installer for Windows does not support custom mode. If you want to override the newer default security settings or upgrade the OpenToken Adapter, use the PingFederate Upgrade Utility.



Important:

If you are upgrading a clustered PingFederate environment, start with the console node, and then follow the additional steps to upgrade the engine nodes.

Steps

- 1. Download the PingFederate installer for Windows from the Ping Identity website.
- 2. Double-click the pingfederate-10.1.0.msi file to begin the upgrade.
- 3. Follow the instructions in the PingFederate Setup Wizard to upgrade PingFederate.

Result:

Errors are rare, but might include problems such as missing or malformed configuration files in the source installation. If the upgrade tool reports an error, review the error messages in the upgrade.log file.



Note:

You can rerun the tool as many times as needed to correct any problems.

When the tool completes the upgrade, it automatically starts the new PingFederate installation.

4. If you are upgrading a clustered PingFederate environment, repeat from step 1 to upgrade PingFederate on each engine node.



Note:

End users might experience disruptions while you upgrade your PingFederate environment.

- 5. Start the new PingFederate installation and open the administrative console.
 - If you are upgrading a clustered PingFederate environment, start the new PingFederate instance on the console node.
- 6. If you are upgrading a clustered PingFederate environment:
 - a. Start the new installation on each engine node, and then ensure all nodes are shown on the System # Server # Cluster Management window.
 - b. Click Replicate Configuration on the Cluster Management window.

Upgrading PingFederate on Windows using the Upgrade Utility

If you used the server distribution . zip file to install the previous version of PingFederate on Windows, you must use the Upgrade Utility to upgrade to the current version of PingFederate.

Before you begin

Read the *Upgrading PingFederate* on page 128 topic for an overview of the upgrade process.

About this task

Upgrade results are contained in the upgrade.log file, which is located in <pf install target>\pingfederate\upgrade\log.



Important:

If you are upgrading a clustered PingFederate environment, start with the console node, and then follow the additional steps to upgrade the engine nodes.

Steps

- 1. Download the latest version of the PingFederate Server distribution .zip file from the Ping Identity website.
- 2. Extract the distribution.zip file into the target installation directory.
- 3. Stop PingFederate.
- 4. At the command prompt, change the current directory to <pf install>\pingfederate\upgrade \bin within the target installation and enter the following command.

```
upgrade <pf install source> [-l <newPingFederateLicense>] [-c]
```

where:

<pf install source>

The full or relative path of the base directory where the existing PingFederate software (pingfederate) is installed.



Note:

The pingfederate subdirectory must exist by that name for the Upgrade Utility to function correctly.

<newPingFederateLicense>

The optional path and file name of the license to use for the upgraded PingFederate version.



Note:

If your current license is valid, the Upgrade Utility automatically copies it from the source installation to the target installation, and you do not need to specify the <newPingFederateLicense> parameter.

If your license is not valid, obtain a valid license file and specify its path and file name for this parameter.

-c

The optional parameter to run the tool in custom mode, which allows you to override newer default security settings (if any) and to upgrade the program files for the OpenToken Adapter.

Result:

The command prompt displays messages indicating upgrade progress. The process is complete when the following message appears.

```
Upgrade completed with [N] errors and [N] warnings
```

If there are errors, scroll up the command window to see them and then correct the indicated problems. Errors during the upgrade should be rare but might include problems such as missing or malformed configuration files in the source installation. The messages are also logged to the upgrade.log file in the Upgrade Utility base directory.

5. If you are upgrading a clustered PingFederate environment, repeat from step 1 to upgrade PingFederate on each engine node.



Note:

End users might experience disruptions while you upgrade your PingFederate environment.

- 6. Start the new PingFederate installation and open the administrative console.
 - If you are upgrading a clustered PingFederate environment, start the new PingFederate instance on the console node.
- 7. If you are upgrading a clustered PingFederate environment:
 - a. Start the new installation on each engine node, and then ensure all nodes are shown on the System # Server # Cluster Management window.
 - b. Click Replicate Configuration on the Cluster Management window.
- 8. If PingFederate is running as a service, re-install the service.
 - a. Remove the existing PingFederate service (see *Uninstalling PingFederate from a Windows server* on page 126).
 - b. Install the new PingFederate service (see Installing PingFederate service on Windows manually on page 126).

For a clustered PingFederate environment, re-install the PingFederate service on all nodes.

- 9. The upgrade utility automatically merges, migrates, and copies the language packs' .properties files into the upgraded PingFederate installation. Verify the language packs in the upgrade installation by looking at the .properties files located in the upgraded cpf install>\pingfederate \server\default\conf\language-packs directory.
 - Standard .properties files include pingfederate-email-messages.properties, pingfederate-messages.properties, and pingfederate-smsmessages.properties. During upgrade, these files are migrated and merged into the upgraded PingFederate installation.
 - Localized .properties files (for example, pingfederate-messages fr CA.properties), are also migrated and merged into the upgraded PingFederate installation.
 - All other .properties files in <pf install>\pinqfederate\server\default\conf \language-packs that do not fit the previous two criteria are copied (not merged) into the upgraded PingFederate installation.

Upgrading PingFederate on Linux systems

On Linux servers, use the Upgrade Utility to upgrade to the current version of PingFederate.

Before you begin

Read the *Upgrading PingFederate* on page 128 topic for an overview of the upgrade process.

About this task

Upgrade results are contained in the upgrade.log file, which is located in <pf install target>/ pingfederate/upgrade/log.



Important:

If you are upgrading a clustered PingFederate environment, start with the console node, and then follow the additional steps to upgrade the engine nodes.

Steps

- 1. Download the latest version of the PingFederate Server distribution.zip file from the Ping Identity website.
- 2. Extract the distribution .zip file into the target installation directory.
- 3. Stop PingFederate.
- 4. On the command line, change the current directory to <pf install>/upgrade/bin within the target installation and execute the following command:

```
./upgrade.sh <pf install source> [-l <newPingFederateLicense>] [-c]
```

where:

```
<pf install source>
```

The full or relative path of the base directory where the existing PingFederate software (pingfederate) is installed.



Note:

The pingfederate subdirectory must exist by that name for the Upgrade Utility to function correctly.

<newPingFederateLicense>

The optional path and file name of the license to use for the upgraded PingFederate version.



Note:

If your current license is valid, the Upgrade Utility automatically copies it from the source installation to the target installation, and you do not need to specify the <newPingFederateLicense> parameter.

If your license is not valid, obtain a valid license file and specify its path and file name for this parameter.

-c

The optional parameter to run the tool in custom mode, which allows you to override newer default security settings (if any) and to upgrade the program files for the OpenToken Adapter.

Result:

The command prompt displays messages indicating upgrade progress. The process is complete when the following message appears.

```
Upgrade completed with [N] errors and [N] warnings
```

If there are errors, scroll up the command window to see them and then correct the indicated problems. Errors during the upgrade should be rare but might include problems such as missing or malformed configuration files in the source installation. The messages are also logged to the upgrade.log file in the Upgrade Utility base directory.

- 5. If you are using AWS CloudHSM version 2.0.x:
 - a. Update the CloudHSM client and the CloudHSM Software Library for Java to version 3.1.1 and restart the client.
 - b. Copy <pf install>/pingfederate/lib-ext/pf-aws-cloud-hsm-wrapper.jar to the JAVA HOME/jre/lib/ext directory.
 - c. Copy all of the files under /opt/cloudhsm/java and /opt/cloudhsm/lib to the JAVA HOME/jre/lib/ext directory.
- 6. If you are upgrading a clustered PingFederate environment, repeat from step 1 to upgrade PingFederate on each engine node.



Note:

End users might experience disruptions while you upgrade your PingFederate environment.

- 7. Start the new PingFederate installation and open the administrative console.
 - If you are upgrading a clustered PingFederate environment, start the new PingFederate instance on the console node.
- 8. If you are upgrading a clustered PingFederate environment:
 - a. Start the new installation on each engine node, and then ensure all nodes are shown on the System # Server # Cluster Management window.
 - b. Click Replicate Configuration on the Cluster Management window.
- 9. If PingFederate is running as a service, re-configure the service.

PingFederate systemd service

Edit the PingFederate systemd unit file and reconfigure the PingFederate service (see step 5 in Installing the PingFederate service on Linux manually on page 123).

PingFederate SysV initialization script

Edit the PingFederate SysV initialization script and reconfigure the PingFederate service (see step 6 in Installing the PingFederate service on Linux manually on page 123).

- 10. The upgrade utility automatically merges, migrates, and copies the language packs' .properties files into the upgraded PingFederate installation. Verify the language packs in the upgrade installation by looking at the .properties files located in the upgraded cpf install>/pingfederate/ server/default/conf/language-packs directory.
 - Standard .properties files include pingfederate-email-messages.properties, pingfederate-messages.properties, and pingfederate-smsmessages.properties. During upgrade, these files are migrated and merged into the upgraded PingFederate installation.
 - Localized .properties files (for example, pingfederate-messages fr CA.properties), are also migrated and merged into the upgraded PingFederate installation.
 - All other .properties files in <pf install>/pingfederate/server/default/conf/ language-packs that do not fit the previous two criteria are copied (not merged) into the upgraded PingFederate installation.

Custom mode

The custom-mode feature in the Upgrade Utility (invoked with the -c option on the command line) allows you to override several default security settings. In addition, if the installed OpenToken Adapter is out of date, running the tool in custom mode allows you to replace the adapter with the latest version.

Security defaults

Using the security defaults should not cause significant issues for most PingFederate installations. The more recent default security settings include:

- Disabling weaker cipher suites for both the SUN and LUNA Java Cryptography Extension (JCE) in PingFederate version 6.2 and later. To see which cipher suites are commented out, choose yes (y) when prompted on whether to use the new defaults. After the upgrade is complete, refer to one of the following configuration files in the new installation's install>/pingfederate/server/ default/data/config-store directory:
 - com.pingidentity.crypto.SunJCEManager.xml
 - com.pingidentity.crypto.AWSCloudHSMJCEManager.xml
 - com.pingidentity.crypto.LunaJCEManager.xml
 - com.pingidentity.crypto.NcipherJCEManager.xml
 - com.pingidentity.crypto.BCFIPSJCEManager.xml

Adapter upgrade

Upgrading the OpenToken Adapter from an earlier version will not normally require any follow-on configuration changes.

- If your existing installation uses a version of the OpenToken Adapter earlier than 2.3, upgrading requires minor configuration modifications in the PingFederate console and redeployment of the agent configuration file.
- If you are upgrading from an OpenToken version earlier than 2.5.1, we recommend that you redeploy agent configuration files, if applicable, as well as any new agent libraries contained in recent versions of PingFederate integration kits and other plug-ins that use OpenToken.



Starting in PingFederate 7.2, the LDAP Java Adapter is no longer supported. This adapter was deprecated in PingFederate 6.6 and replaced by the LDAP Username Password Credential Validator (PCV), which can be used with the HTML Form Adapter or HTTP Basic Adapter.

Reviewing post-upgrade tasks

Review the following topics for post-upgrade tasks:

- Reviewing administrative users on page 143
- Copying customized files or settings on page 144
- Reviewing database changes on page 146
- Reviewing log configuration on page 151
- Migrating other components on page 152
- Resetting files and variable for HSM on page 154
- Verifying the new installation on page 154

It is also important to perform runtime tests to ensure the new PingFederate installation fulfills your existing use cases.

Reviewing administrative users

As of PingFederate 10.1, the use of expressions is enabled by default. Additionally, a new administrative role, Expression Admin, has been added.

When upgrading to PingFederate 10.1 from a previous version, administrative users who were granted the Admin role in the earlier installation are granted the Expression Admin role automatically. You can achieve the same result by using the /bulk/import administrative API endpoint to bulk-import a configuration that was bulk-exported from PingFederate 10.0.

If preferred, administrators can disable the use of expressions by setting evaluateExpressions to false as described in Enabling and disabling expressions on page 215. Also, go to the System # Server # Administrative Accounts window and remove the Expression Admin role from all Admin users. Doing this will prevent Admin users from entering expressions into PingFederate if the evaluateExpressions element is set to true at a later time. For more information, see *Administrative accounts* on page 842.

Copying customized files or settings

After you have upgraded PingFederate, you must copy files that were customized in the previous release to the current installation.

User-facing windows

If you have modified any Velocity templates for user-facing windows, to preserve the customized user experience, you must migrate your custom changes to the new installation manually for each server node. The templates are located in the <pf install>/pingfederate/server/default/conf/template directory.



CAUTION:

Supporting CSS and image file names were changed as of PingFederate 7.0. For each modified HTML template copied, add .1 to the base name for each CSS file referenced in the header (for example <link rel="..." href="assets/css/window.1.css"/>).

Add .1 to any references in the copied templates to the installed image files contained in the assets/ images directory (for example, <imq src="assets/images/green check.1.png"/>)

Email notifications

If you modified the email notification templates prior to version 9.2, manually migrate your custom changes to the new HTML-based templates for each server node.

The plain text templates (message-template-*.txt) are located in <pf install>/pingfederate/ server/default/conf in the source installation. The new HTML-based templates are located in <pf install>/pingfederate/server/default/conf/template/mail-notifications with the same file naming convention but an .html file extension.

Jetty or JBoss configuration

If you have modified any Jetty or JBoss settings that need to be carried forward, you must make the corresponding changes manually in the new PingFederate deployment.

If you are upgrading from version 6.9 or later, you can copy over the relevant files from the Jetty configuration directory, <pf install>/pingfederate/etc.

If you are upgrading from version 6.0 through 6.8, first identify any changes made to the JBoss configuration, then make corresponding changes for the newer Jetty configuration.

For example, if you modified the <pf install>/pingfederate/server/default/deploy/ jetty.sar/META-INF/jboss-service.xml file prior to version 6.9, identify the changes and make the same modifications at corresponding points in either the jetty-admin.xml or jetty-runtime.xml files located in the new Jetty configuration directory, <pf install>/pingfederate/etc.

size-limits.conf

Prior to version 8.4.2, the InterReqStateMgmtMapImpl.expiry.mins Setting in the cpf install>/ pingfederate/server/default/conf/size-limits.conf file defines the lifetime of the following data sets:

Inter-request state information

The state information between the redirects to complete a request.

Adapter session-state data

The state information (along with the associated attributes and their values, if any) maintained (or used) by the adapters.

PingFederate 8.4.2 splits the InterReqStateMgmtMapImpl.expiry.mins settings into two settings, one setting for each data type.

New settings	Data type	Default value in minutes
InterReqStateMgmtMapImpl.expiry.mins.sta	antermanuest state information	30
InterReqStateMgmtMapImpl.expiry.mins.at	Adapter session-state data	1440 (24 hours)

The new settings reduce the memory footprint of PingFederate by purging the inter-request state information after 30 minutes and retaining adapter session-state data during the day.

If you have previously modified the value of the InterReqStateMgmtMapImpl.expiry.mins setting, when migrating your change to the latest version, adjust the value of the new settings based on your requirements.

Cross-origin resource sharing (CORS) support for OAuth endpoints

If you have previously edited the <pf install>/pingfederate/etc/webdefault.xml file to enable CORS support for OAuth endpoints, instead of updating the webdefault.xml file, define the allowed origins manually using the PingFederate administrative console after the upgrade. For more information, see Configuring authorization server settings on page 475.

Configuration files in the config-store directory

If you have added or replaced setting values in configuration files stored in the <pf install>/ pingfederate/server/default/data/config-store directory, the PingFederate upgrade tools copy these setting values to the new installation.



The upgrade tools do not copy comments from the existing installation to the new installation.

If you have removed a setting or a block of settings from a configuration file in the config-store directory, the upgrade tool preserves your changes by removing the setting or block of settings from the new installation and records the removals in its log file. To re-add a setting or block of settings to the new installation, compare the configuration file found in the new installation to the file found in the product distribution . zip file and make your changes.

Other configuration files

As of PingFederate 10.0, the upgrade process copies many files automatically. However, there are still some files that you must copy manually. All of these files are located in <pf install>/pingfederate/ server/default/conf.

The following files are copied automatically:

- Properties files with the .conf extension
- The log4j2.db.properties file
- The jmx-remote-config.xml file
- Non-default files located in the template directory

If you have modified the default templates (located in <pf install>/pingfederate/server/ default/conf/template, you must customize these templates in the new PingFederate installation.

If you modified versions of tcp.xml, udp.xml, and log4j2.xml, they are copied over intact. The default files are saved in the target directory with a different extension. To take advantage of the improvements in the default versions of these files, migrate your changes into the current default files and then rename them appropriately.



If you are upgrading from version 8.0 or earlier, PingFederate might not start until you have merged your changes into the current default files due to JGroups errors.

Other files, such as jmx.remote.access, are not copied to the new installation automatically. To preserve any custom settings, create a backup of the current configuration files and merge your changes to the current files.

If you have previously customized Java Virtual Machines (JVM) options in the run.bat or run.sh files, instead of updating these files, manually merge your JVM options to the <pf install>/ pingfederate/bin/jvm-memory.options file. For more information, see Fine-tuning JVM options on page 1003 and *memoryoptions and upgrade* on page 999.

Reviewing database changes

Occasionally, PingFederate introduces database-related changes, such as adding a new table, modifying an existing table, or updating connection pool library, for the purpose of product improvement.

Neither the Upgrade Utility nor the PingFederate installer for Windows migrates data maintained in the internal HSQLDB database or any external database. For instance, if outbound provisioning is enabled in the new PingFederate instance using the internal database, it is re-initialized from the provisioning source.



CAUTION:

Use the built-in HSQLDB only for trial or training environments. For testing and production environments, always use a secured external storage solution for proper functioning in a clustered environment.

Testing involving HSQLDB is not a valid test. In both testing and production, it might cause various problems due to its limitations and HSQLDB involved cases are not supported by Pingldentity.

If your PingFederate environment connects to one or more database servers, review the following topics and make changes accordingly.

- Provisioning datastore reset on page 147
- Security enhancement in JDBC datastore gueries on page 147
- New connection pool library on page 148
- An improved index in the database table for OAuth clients on page 149

- Changes in the database tables for log messages on page 149
- Changes in the database table for account linking on page 149
- Changes in the database tables for OAuth clients on page 149
- Changes in the database tables for OAuth persistent grants and extended attributes on page 149
- A new database table for OAuth persistent grant extended attributes on page 150
- New indexes in the database table for OAuth persistent grants on page 150
- Changes in a database table supporting nested group membership on page 150

Provisioning datastore reset

About this task

Upgrading to PingFederate 9.0 or 9.0.1 when using its outbound provisioning capability can result in user records being disabled at SaaS applications. The issue has since been resolved in version 9.0.2.

If you are upgrading from version 8.4.4 (or an earlier version) or 9.0.2 to version 10.1, the upgrade process automatically resolves this issue. No further action is required.

If you are upgrading from version 9.0 or 9.0.1 to version 10.1, use the following provmgr commands to reset the provisioning datastore on the upgraded installation:

The provmgr command-line tool, used in the following steps, is located in the <pf install>/ pingfederate/bin directory: provmgr.bat for Windows and provmgr.sh for Linux. For more information about the provmgr command-line tool, see Outbound provisioning CLI on page 795.

Steps

1. Run the following command to obtain a list of provisioning channel IDs.

```
provmgr --show-channels
```

2. Reset the provisioning datastore for a given channel by its ID.

```
provmgr -c <channel id> --reset-all
```



Note:

If you have multiple provisioning channels, run the command for each channel. The order of the parameters does not matter.

Security enhancement in JDBC datastore gueries

A security enhancement has been made in PingFederate 9.0 to safeguard JDBC datastore queries against back-end SQL injection attacks. This protection is enabled for all new installations.

About this task

For upgrades, you can enable this protection by modifying the f install>/pingfederate/ server/default/data/config-store/org.sourceid.common.SqlFilterManager.xml file.

To enable this security enhancement:

Steps

- 1. Edit the org.sourceid.common.SqlFilterManager.xml file.
- 2. Set the <item name="enableSqlFilters"/> element value to true; for example:

```
<?xml version="1.0" encoding="UTF-8"?>
<config xmlns="http://www.sourceid.org/2004/05/config">
```

```
<item name="enableSqlFilters">true</item>
</config>
```

- 3. Save the file.
- 4. Restart PingFederate.

If you have a clustered PingFederate environment:

- a. Perform the previous steps on the console node.
- b. Sign on to the PingFederate administrative console.
- c. Go to the System # Server # Cluster Management window.
- d. Click Replicate Configuration to push this change to all engine nodes.
- 5. Verify your use cases to make sure your search filters return the expected results.

New connection pool library

As of PingFederate 8.0, support for BoneCP as the JDBC connection pool library has been deprecated and replaced with Apache Commons DBCP 2, which requires JDBC 4.1 (or more recent) drivers.

About this task

Verify that the database-driver .jar files, located in the <pf install>/pingfederate/server/ default/lib directory, meet the minimum version requirement. If you are using JDBC drivers of version 4.0 (or earlier), contact your vendors for the latest drivers and replace the older JDBC database-driver .jar files with the latest.



Important:

Support for BoneCP as the JDBC connection pool library has been deprecated and will be removed in a future release.



Note:

PingFederate was tested with vendor-specific JDBC drivers. For more information, see *Database driver* information on page 115. To obtain the database driver JAR file, contact your database vendor. Install the database driver file to the <pf install>/pingfederate/server/default/lib directory then restart the server.

To upgrade the older JDBC drivers at a later time, use the following steps to revert the JDBC connection pool library to BoneCP.

Steps

- 1. Edit the <pf install>/pingfederate/server/default/data/config-store/ com.pingidentity.jdbc.DataSourceDeployerFactory.xml file.
- 2. Replace dbcp with bonecp.
- 3. Save the file.
- 4. Restart PingFederate.

If you have a clustered PingFederate environment:

- a. Perform the previous steps on the console node.
- b. Sign on to the PingFederate administrative console.
- c. Go to the **System # Server # Cluster Management** window.
- d. Click Replicate Configuration to push this change to all engine nodes.

An improved index in the database table for OAuth clients

PingFederate 8.4 added the value column to an existing index (IDX_FIELD_NAME) in the pingfederate_oauth_clients_ext table as a general improvement.

This information is applicable only to customers who configured PingFederate to store OAuth clients on a database server.

You must modify the index in your existing pingfederate_oauth_clients_ext table.

While there is no alter-table script provided, you can derive the setup from the new table-setup scripts, oauth-client-management-<databaseServer>.sql, found in the <pf_install>/pingfederate/server/default/conf/oauth-client-management/sql-scripts directory.

Changes in the database tables for log messages

For performance reasons, PingFederate 8.0 started using the NVARCHAR data type instead of the VARCHAR data type.

This information is applicable only to customers who have configured Log4j or Log4j 2 to write log messages to database tables on Microsoft SQL Server.

The table-setup scripts (targeted for Microsoft SQL Server) in the $<\!pf_install>/pingfederate/$ server/default/conf/log4j/sql-scripts directory have been updated. If you are upgrading from PingFederate 7.3 or an earlier version, consider updating the data type in the applicable tables accordingly for performance reasons.

Changes in the database table for account linking

For columns in the pingfederate_account_link table using the VARCHAR data type, PingFederate 7.3 updated the data type to NVARCHAR for performance reasons.

This is information is applicable only to customers who have created a database table for account linking on Microsoft SQL Server.

The table-setup script, $<\!pf_install>\!/pingfederate/server/default/conf/account-linking/sql-scripts/account-linking-sqlserver.sql, has been updated. If you are upgrading from PingFederate 7.2 R2 or an earlier version, consider updating the data type in the pingfederate_account_link table accordingly for performance reasons. Consult with your database administrator, as needed.$

Changes in the database tables for OAuth clients

For columns in the pingfederate_oauth_clients and pingfederate_oauth_clients_ext tables using the VARCHAR data type, PingFederate 7.3 updated the data type to NVARCHAR for performance reasons.

This information is applicable only to customers who have deployed Microsoft SQL Server to store their OAuth clients.

The table-setup script, $<\!pf_install>\!$ /pingfederate/server/default/conf/oauth-client-management/sql-scripts/oauth-client-management-sqlserver, has been updated. If you are upgrading from PingFederate 7.2 R2 or an earlier version with OAuth use cases, consider updating the data type in both tables accordingly for performance reasons. Consult with your database administrator, as needed.

Changes in the database tables for OAuth persistent grants and extended attributes

For columns in the pingfederate_access_grant and pingfederate_access_grant_attr tables using the VARCHAR data type, PingFederate 7.3 updated the data type to NVARCHAR for performance reasons.

This information is applicable only to customers who have deployed Microsoft SQL Server to host their OAuth grant datastore.

The table-setup scripts (targeted for Microsoft SQL Server) in the cpf_install>/pingfederate/server/default/conf/access-grant/sql-scripts directory have been updated. If you are upgrading from PingFederate 7.2 R2 or an earlier version with OAuth use cases, consider updating the

data type in both tables accordingly to avoid potential issues caused by incompatible data types between these two tables. Consult with your database administrator, as needed.

A new database table for OAuth persistent grant extended attributes

PingFederate 7.2 R2 introduced the capability to map attributes from initial authentication context to an access token attribute contract, which requires an update in the OAuth grant datastore.

If you are upgrading from PingFederate 7.2.1 or an earlier version with OAuth use cases, run the tablesetup script, <pf install>/pingfederate/server/default/conf/access-grant/sglscripts/access-grant-attribute-<database>.sql, for your database server.

New indexes in the database table for OAuth persistent grants

Since PingFederate 7.1 R3, two indexes have been added to the pingfederate_access_grant table.

PingFederate version	Column	Index
7.3	expires	EXPIRESIDX
7.1 R3	client_id	CLIENTIDIDX

If you are upgrading from PingFederate 6.5 through 7.1.4, you must add both indexes to your existing pingfederate_access_grant table.

If you are upgrading from PingFederate 7.1 R3, 7.2.x, or 7.2 R2, you must add the EXPIRESIDX index for the expires column.

While there is no alter-table script provided, you can derive the setup from the new table-setup scripts, <pf install >/pingfederate/server/default/conf/access-grant/sql-scripts/ access-grant-<databaseServer>.sql.

Changes in a database table supporting nested group membership

Outbound provisioning of groups and nested group membership requires an update in the internal datastore.

If you are upgrading from PingFederate 8.0.x, no action is required.

If you are upgrading from PingFederate 7.1.x through 7.3, consider running the alter-table script, <pf install>/pingfederate/server/default/conf/provisioner/sql-scripts/updates/ add-subgroup-membership-<database>.sql, for your database server to update the existing group_membership table.

If you are upgrading from PingFederate 7.0.1 or an earlier version, consider using the tablesetup script, <pf install>/pingfederate/server/default/conf/provisioner/sqlscripts/provisioner-<database>.sql, for your database server as a reference to add the new group_membership table to your existing internal datastore.

Alternatively, you can create a new database using the table-setup script and let the outbound provisioner repopulate the new internal datastore on its next run, regardless of the PingFederate version you are running.



CAUTION:

The table-setup script removes the existing (outbound provisioning) tables. To keep a copy of the current data, use the tools available from your database vendor to make a backup before running the table-setup script.

Reviewing log configuration

Starting with version 8.0, PingFederate uses Log4j2 to write log messages. This upgrade improves performance and allows real-time log level adjustments. If you have not modified the Log4j configuration file (log4j.xml) in an earlier version, PingFederate 8.0 starts using Log4j2 after the upgrade process.

If the configuration file for Log4j 2 (or Log4j) has been modified in the source installation, refer to one of the following sections for upgrade instructions.

Upgrading from PingFederate 8.x, 9.x, or 10.x

About this task

If the Upgrade Utility or the PingFederate installer for Windows determines that the Log4j2 configuration file ($<pf_install>/pingfederate/server/default/conf/log4j2.xml)$ has changed since it was originally installed, new features are not activated; the upgrade tools do not support automatic merging of customizations made to the existing logging configuration. Instead, these upgrade tools copy the modified log4j2.xml file to the new installation intact and rename the configuration file from the product.zip file using the new PingFederate version number. Both configuration files are located in the same conf directory.

To activate new features:

Steps

- 1. Review the new features by comparing the renamed Log4j2 configuration file against log4j2.xml.
- 2. Modify log4j2.xml to suit your needs.
- 3. If you have a clustered PingFederate environment, repeat step 2 for all applicable PingFederate nodes in the cluster.

Upgrading from PingFederate 6.x or 7.x

About this task

PingFederate 6.x and 7.x use Log4j to write log messages.

If the Upgrade Utility or the PingFederate installer for Windows determines that the Log4j configuration file ($<pf_install>/pingfederate/server/default/conf/log4j.xml$) has changed since it was originally installed, these upgrade tools copy the modified log4j.xml to the new installation with a new name and installs the new log4j2.xml from the product.zip file.

The new name for the previously customized log4j.xml is log4j-old-<SourceVersion>.xml, where <SourceVersion> is the version number of the source PingFederate installation.

Both configuration files are located in the conf directory.

To migrate custom changes from Log4j to Log4j2:

Steps

1. Review custom changes by comparing log4j-old-<SourceVersion>.xml the pf-upgrade<NewVersion>/reference-files/<SourceVersion>/server/default/conf/log4j.xml file from the upgrade tool.



Important:

The <SourceVersion> values must match.

2. Modify the log4j2.xml file to suit your needs.



The configuration syntax between Log4j and Log4j2 varies. For more information, see the Apache Log4j 2 documentation.

- 3. If you are writing log messages to a database server, enter the database information into log4j2.db.properties in the same conf directory.
- 4. If you have a clustered PingFederate environment, repeat steps 2 and 3 for all applicable PingFederate nodes in the cluster.

Migrating other components

Some custom and integrated components might require additional steps after upgrading PingFederate.

Updating the custom authentication selector

Through the use of the PingFederate SDK, you can create a custom authentication selector by implementing the AuthenticationSelector interface. Most implementations return AuthenticationSelectorContext.ResultType.CONTEXT as the result type, which requires no further action after an upgrade.

If your implementation returns AuthenticationSelectorContext.ResultType.ADAPTER ID (an IdP adapter instance ID) or AuthenticationSelectorContext.ResultType.IDP CONN ID (the connection ID of an IdP connection), you must update the descriptor instance of your custom authentication selector to call the setSelectAuthnSourceResultType method with an input of true. For each authentication policy path that ends with an instance of such custom authentication selector, you must ensure that its action is set to Done.

For more information, refer to the Javadoc for the AuthenticationSelector interface and the AuthenticationSelectorDescriptor class.



The Javadoc for PingFederate is located in the <pf install>/pingfederate/sdk/doc directory.

Migrating to the integrated LDAP Username PCV

About this task

As of PingFederate 7.3, the integrated LDAP Username Password Credential Validator (PCV) can return additional attribute values upon successful validation. If you have previously deployed the LDAPExtendedAttributesPCV-<version>.jar file from the PingID integration kit and created an instance of the LDAP PCV with Extended Attributes, migrate to the integrated LDAP Username PCV.

Steps

- 1. Create an instance of the integrated LDAP Username PCV.
 - a. On the System # Data & Credential Stores # Password Credential Validators window, click Create New Instance.
 - b. On the Type tab, enter the required information and select LDAP Username Password Credential Validator from the list.
 - c. On the Instance Configuration tab, select an LDAP datastore from the list, enter a search base and a search filter, and select the scope of the search.



You can reuse the information from the existing LDAP PCV with Extended Attributes instance.

- d. On the Extended Contract tab, enter memberOf under Extend the Contract and click Add.
- e. On the **Summary** tab, review the setup and click **Done**.
- f. On the Manage Credential Validator Instances window, click Save.
- 2. In the configuration where the LDAP PCV with Extended Attributes instance is used, replace it with the newly created LDAP Username Password Credential Validator instance.

Example:

For example, if you have created an instance of the PingID PCV (with integrated RADIUS server) instance and have selected an instance of the LDAP PCV with Extended Attributes as one of the delegate PCVs, remove the selection and add the newly created LDAP Username Password Credential Validator instance to the list.

- 3. After replacing the LDAP PCV with Extended Attributes instance, delete it from the Password Credential Validators window.
- 4. Remove the <pf install>/pingfederate/server/default/deploy/ LDAPExtendedAttributesPCV-<version>.jar file on all PingFederate servers.
- 5. Restart PingFederate on all PingFederate servers.

Migrating to the integrated Username Token Processor

About this task

As of PingFederate 7.2, the Username Token Translator has been deprecated and replaced with an integrated Username Token Processor. While the integrated Username Token Processor and the deprecated Username Token Translator may be simultaneously deployed, it is recommended to migrate to the new token processor.

Steps

- 1. Go to the **Identity Provider # Token Processors** screen.
- 2. Click Create New Instance to create an instance of the integrated Username Token Processor.



Important:

In the **Type** screen, select **Username Token Processor** from the list.



If you have multiple WS-Trust STS SP connections, you may reuse the same Username Token Processor instance or create additional instances of the token processors as needed.

- 3. Map the new token processor instance to the applicable WS-Trust STS SP connection on the IdP Token Processor Mapping screen.
 - Repeat this step if you have multiple WS-Trust STS SP connections.
- 4. Test your WS-Trust STS SP connections using the instance of the integrated Username Token Processor.
- 5. Remove the token processor instance of the deprecated Username Token Translator from all WS-Trust STS SP connections in the IdP Token Processor Mapping screen.
- 6. If you have set up token translator mappings, create new entries to replace those using instances of the deprecated Username Token Translator, test the new mapping entries, and delete the entries that use instances of the deprecated Username Token Translator.
- 7. Delete all token processor instances of the deprecated Username Token Translator in the **Identity** Provider # Token Processors screen.
- 8. Remove the pf-username-token-translator-<version>.jar file from the <pf install>/ pingfederate/server/default/deploy directory on all PingFederate servers.
- 9. Restart PingFederate on all PingFederate servers.

Resetting files and variable for HSM

If your PingFederate installation is configured in a clustered environment with nCipher nShield Connect, you must copy the <pf install>server/default/data/ncipher-kmdata-local directory to the new installation manually and update the environmental variable NFAST KMLOCAL to point to the new location.

Verifying the new installation

Integration kits and custom solutions might introduce third-party dependencies that could trigger runtime errors. We recommend that you perform runtime tests, verifying that the previously deployed use cases, including any OGNL expressions, are functional in the new installation.

Getting Started with PingFederate Server 10.1

This guide provides information about configuring PingFederate to deploy a secure Internet-identity platform, including single sign-on (SSO), based on the latest security and business standards.

Starting and stopping PingFederate servers

Depending on the application mode and the operating system, the steps to start, stop, or restart PingFederate vary.

When you install or upgrade PingFederate using its platform-specific installer, PingFederate configures to run as a service. You can stop and disable the service and run PingFederate as a console application.

If you install or upgrade PingFederate manually by using the PingFederate product distribution file or the Upgrade Utility in command line, you can run PingFederate as a console application or install the PingFederate service manually and run it as a service.

Starting and stopping PingFederate on Windows

Steps

• Follow the relevant steps to start PingFederate on Windows.

Application mode	Steps
Console application	 Open a command prompt. Go to <pf_install>/pingfederate/bin.</pf_install> Run run.bat. Keep the command prompt open.
Windows service	Go to Control Panel # System and Security # Administrative Tools # Services. Right-click on the PingFederate service and click Start.

Follow the relevant steps to stop PingFederate on Windows.

Application mode	Steps
Console application	Locate the command promptrunning PingFederate. Press CTRL+C to terminate PingFederate.
Windows service	 Go to Control Panel # Administrative Tools # Services. Right-click on the PingFederate service and click Stop.

• Follow the relevant steps to restart PingFederate on Windows.

Application mode	Steps
Console application	 Locate the command prompt running PingFederate. Press CTRL+C to terminate PingFederate. When PingFederate stops, run run.bat. Keep the command prompt open.
Windows service	 Go to Control Panel # Administrative Tools # Services. Right-click on the PingFederate service and select Restart.

Starting and stopping PingFederate on Linux

Steps

Follow the relevant steps to start PingFederate on Linux.

Application mode	Steps
Console application	Open a terminal window.

Application mode	Steps
	 Go to <pf_install>/pingfederate/bin.</pf_install> Run run.sh. Keep the terminal window open.
Service	 Open a terminal window. Enter the system-dependent service command to start PingFederate.

Follow the relevant steps to stop PingFederate on Linux.

Application mode	Steps
Console application	Locate the terminal window running PingFederate. Press CTRL+C to terminate PingFederate.
Service	Open a terminal window. Enter the system-dependent service command to stop PingFederate.

Follow the relevant steps to restart PingFederate on Linux.

Application mode	Steps
Console application	 Locate the terminal window that is running PingFederate. Press CTRL+C to terminate PingFederate. When PingFederate stops, run run.sh. Keep the terminal window open.
Service	Open a terminal window. Enter the system-dependent service command to restart PingFederate.

Opening the PingFederate administrative console

The PingFederate administrative console provides a wizard-like interface in which you configure your federation use cases.

About this task

To open the administrative console:

Steps

- 1. Start PingFederate. See Starting and stopping PingFederate servers on page 154. In a clustered PingFederate environment, start PingFederate on the console node.
- 2. Start a web browser.

3. Go to https://*<pf host>*:9999/pingfederate/app.



<pf_host> is the network address of your PingFederate server. It can be an IP address, a host name, or a fully qualified domain name. It must be reachable from your computer.

9999 is the default value of the pf.admin.https.port property in the run.properties file.

Setting up PingFederate

The first time you access the PingFederate administrative console and after you have accepted the subscription agreement, you see two choices on how to proceed:

- Yes, Connect to PingOne for Enterprise
- No, Set Up Without PingOne for Enterprise

To continue setting up PingFederate, select No, Set Up Without PingOne for Enterprise. Click Next.



Important:

Selecting Yes, Connect to PingOne for Enterprise takes you to the setup for PingFederate Bridge. PingFederate Bridge is a light-weight version of PingFederate that is meant for use primarily by new customers who want to quickly and easily configure user authentication from an on-premise directory to PingOne for Enterprise in the cloud. For information about setting up and using PingFederate Bridge, see the associated documentation starting with Introduction to PingFederate Bridge.

Importing your license

In the setup wizard, on the License tab, import your license.

Before you begin

Make sure you have a license or request a license key at the Ping Identity licensing website or contact sales@pingidentity.com.

About this task

If a license is pre-installed in your folder, the License Summary displays the relevant license information.

Steps

- 1. To locate and import your license file, click Choose File.
- 2. Review your License Summary. Click Next.

You can replace the current license with another license file later, see *Installing a replacement license* key.

Entering basic information for PingFederate

On the **Basic Information** tab, enter your federation information.

Steps

1. In the **Base URL** field, verify your base URL. Update as needed.

The domain portion of the base URL should match the domain name of your organization because it is part of the address where your applications, users, and partners communicate with your PingFederate environment.



You can add multiple virtual host names at a later time. For more information, see Virtual host names on page 854.

2. In the **Entity ID** field, enter your Entity ID if prompted. Click **Next**.



Note:

This is the unique identifier of your organization. It is how your partners identify you when communicating with you based on SAML 2.0 specifications.

Configuring identity provider settings

If you have selected the Identity Provider role on the Enable Roles tab, the Identity Provider Configuration tab becomes available.

About this task

On the **Identity Provider Configuration** tab, connect PingFederate to a directory server.



If the directory is Microsoft Active Directory, you can enable Kerberos authentication for Windows users.

Steps

- 1. Click **Begin** and then follow the instructions to complete the configuration. Result: The administrative console returns you to the **Identity Provider Configuration** tab.
- 2. Click Next.

Connecting PingFederate to a directory

On the **Connection** tab, connect PingFederate to a directory.

Steps

From the **Directory Type** list, select a directory type and provide the required information. Click **Next**. For information about each field, see the following table.

Field	Description
Directory Type	Select the directory server type from the list.
	For a list of supported directory servers, see <i>System requirements</i> on page 110.

Field	Description
Data Store Name	Enter the data store name.
Hostname	Enter directory server location.
	It can be the IP address, the host name, or the fully-qualified domain name of the directory server. The entry can include a port number.
Service Account DN	Enter the distinguished name (DN) of the service account that PingFederate can use to communicate with the directory server.
Password	Enter the service account password.
Search Base	Enter the DN of the location in the directory where PingFederate begins its data store queries.
Search Filter	Enter the LDAP query to locate a user record for attribute lookup and potentially credential validation.
	Depending on the selected directory type, the default value is either sAMAccountName=\${username} or uid=\${username}.
	If you require a more advanced search filter, ensure the value is a valid LDAP filter. For more information, consult your directory administrators.

Result:

When you click **Next**, PingFederate tries to establish a secure LDAPS connection to the directory server.

If the directory server does not support LDAPS, the Unsecure Connection window appears. To continue without a secure connection, click Next. Alternatively, you can go back to the Connection tab and enter a different directory server.

If PingFederate does not trust the certificate from the directory server, the Certificate Error window appears. Import the directory server certificate and then click Next. Alternatively, you can go back to the **Connection** tab and enter a different directory server.

Configuring Kerberos authentication

If you have chosen to connect PingFederate to Microsoft Active Directory, the Kerberos Authentication tab becomes available.

About this task

Enable Kerberos authentication for Windows users on the Kerberos Authentication tab.



Important:

If you have not configured a service account for Kerberos authentication, see Configuring the Active Directory environment on page 916.

Steps

Select the Configure Kerberos Authentication check box and provide the required information. Click Next.

For information about each field, refer to the following table.

Field	Description
Realm Name	Enter the fully-qualified domain name.

Field	Description
Realm Username	Enter the service account that PingFederate can use to communicate with Active Directory for the purpose of Kerberos authentication.
Realm Password	Enter the service account password.
Internal IP Ranges	Enter one or more network ranges where PingFederate can try authenticating with the Kerberos protocol when handling requests originating from such IP addresses.
	Typically, these are internal network ranges with access to one or more key distribution centers (KDCs) in your domain.
	To remove an entry, select it from the list and then click Delete .
KDC Hostnames	Enter the host name or the IP address of the applicable KDC.
(Optional)	This field is optional. Multiple hosts are allowed. If left unspecified, PingFederate uses a DNS query to find a list of KDCs.
	To remove an entry, select it from the list and then click Delete .

Next steps

Kerberos authentication also requires browser-specific configuration. For more information, see Configuring end-user browsers on page 316.

Reviewing the directory configuration

On the **Summary** tab, review your configuration.

Steps

Perform the following actions as needed.

Action	How to accomplish it
	Click the corresponding tab and follow the configuration workflow
Keep your changes	Click Done and continue with the configuration
Discard your changes	Click Cancel

Creating an administrator account

On the Administrator Account tab, create an administrative account.

Steps

1. Replace the default value in the **Username** field with a username of your choice.

The default value is Administrator.

2. Enter a password in the Password and Confirm Password fields, and then click Next.

Reviewing the PingFederate configuration

On the **Confirmation** tab, review your configuration.

Steps

Perform one of the following actions.

Choose from:

- To amend your configuration, click the corresponding tab, and then follow the configuration wizard.
- To keep your configuration, click Next, and then click Done.
- To discard your changes, close your browser and restart PingFederate.

Setting up PingFederate Bridge

The first time you open the PingFederate Bridge administrative console, you see the initial setup screen, from which you can begin setting up PingFederate Bridge. Each screen has online help, which you can access by clicking the question mark icon (?) located in the top right corner.

Connecting PingFederate to PingOne for Enterprise

Integrating PingFederate with PingOne for Enterprise provides a powerful solution combining the benefits of an on-premise deployment with the flexibility of a cloud solution.

Steps

- 1. Select Yes, Connect to PingOne for Enterprise, and then click Sign on to PingOne to get your activation key.
- 2. Sign on using your PingOne admin portal credentials.



Z Tip:

If you do not have a PingOne for Enterprise account, you can register for a free trial.

- 3. In the PingOne admin portal, copy the value from the **Activation Key** field.
- 4. On the PingOne Account tab of the PingFederate administrative console, paste the value in the Activation Key field. Click Next.



Note:

To connect PingFederate to PingOne for Enterprise later, select No, Set Up Without PingOne for Enterprise on the PingOne Account tab. When you are ready to connect to PingOne for Enterprise, go to System # External Systems # Connect to PingOne for Enterprise.

Connecting PingFederate Bridge to a directory server

PingFederate can use a directory server for PingOne SSO and PingID VPN integration.

Steps

 To connect to a directory, select Yes, Connect a Directory Server on the Identities window of the setup wizard. Provide the required information, and then click Next.

For information about each field, refer to the following table.

Field	Description
Directory Type	Select the directory server type from the list.
	For a list of supported directory servers, see <i>System requirements</i> on page 110.
Data Store Name	Enter the datastore name.

Field	Description
Hostname	Enter the directory server location.
	It can be the IP address, the host name, or the fully qualified domain name of the directory server. The entry may include a port number.
Service Account DN	Enter the distinguished name (DN) of the service account that PingFederate can use to communicate with the directory server.
Password	Enter the service account password.
Search Base	Enter the DN of the location in the directory where PingFederate begins its datastore queries.
Search Filter	Enter the LDAP query to locate a user record for attribute lookup and potentially credential validation.
	Depending on the selected directory type, the default value is either sAMAccountName=\${username} or uid=\${username}.
	If you require a more advanced search filter, ensure the value is a valid LDAP filter. For more information, consult your directory administrators.

Result:

When you click Next, PingFederate tries to establish a secure LDAPS connection to the directory server.

If the directory server does not support LDAPS, the Unsecure Connection window appears. To continue without a secure connection, click Next. Alternatively, you have the option to go back to the **Identities** tab and enter a different directory server.

If PingFederate does not trust the certificate from the directory server, the **Certificate Error** window appears. Import the directory server certificate and then click Next. Alternatively, you have the option to go back to the **Identities** tab and enter a different directory server.

To connect to a directory later, select No, Don't Connect a Directory Server, and then click Next.



This setup scenario is suitable for proof of concept. This scenario creates local test accounts as a result.

Configuring PingOne SSO and PingID VPN settings

Steps

1. To enable PingOne for Enterprise integration, select the PingOne SSO check box.



Note:

If you have chosen to connect to a directory server, select the Additional SSO Features check box to enable provisioning. If the directory is Microsoft Active Directory, selecting this check box allows you to enable Kerberos authentication.

Click Begin.

- 2. To enable and configure a built-in RADIUS server to integrate PingID with your VPN, select the PingID VPN (RADIUS) check box, and then click Begin.
- 3. When the administrative console returns you to this window, click Next.

Configuring Kerberos authentication

If you have chosen to connect PingFederate to Microsoft Active Directory, the Kerberos Authentication tab becomes available.

About this task

Enable Kerberos authentication for Windows users on the Kerberos Authentication tab.



Important:

If you have not configured a service account for Kerberos authentication, see Configuring the Active Directory environment on page 916.

Steps

Select the Configure Kerberos Authentication check box and provide the required information. Click Next.

For information about each field, refer to the following table.

Field	Description
Realm Name	Enter the fully-qualified domain name.
Realm Username	Enter the service account that PingFederate can use to communicate with Active Directory for the purpose of Kerberos authentication.
Realm Password	Enter the service account password.
Internal IP Ranges	Enter one or more network ranges where PingFederate can try authenticating with the Kerberos protocol when handling requests originating from such IP addresses.
	Typically, these are internal network ranges with access to one or more key distribution centers (KDCs) in your domain.
	To remove an entry, select it from the list and then click Delete .
KDC Hostnames	Enter the host name or the IP address of the applicable KDC.
(Optional)	This field is optional. Multiple hosts are allowed. If left unspecified, PingFederate uses a DNS query to find a list of KDCs.
	To remove an entry, select it from the list and then click Delete .

Next steps

Kerberos authentication also requires browser-specific configuration. For more information, see Configuring end-user browsers on page 316.

Configuring provisioning to PingOne for Enterprise

If you are connecting PingFederate to a directory server, the **Provisioning** tab becomes available.

About this task

On the Provisioning tab, you enable provisioning of users and groups from your directory server to PingOne for Enterprise by specifying the group where PingFederate should look for member users.

Steps

1. Select the **Configure Provisioning** check box.

2. In the **Group DN** field, enter the distinguished name (DN) of the applicable group.



The specified group must reside under the hierarchy of the previously-defined Search Base value.

- 3. Optional: For PingFederate to provision users through nested group membership, select the Nested check box.
- 4. Click Next.

Reviewing the PingOne SSO configuration

On the **Summary** tab, review your configuration.

Steps

Perform the following actions as needed.

Action	How to accomplish it
	Click the corresponding tab and follow the configuration workflow
Keep your changes	Click Done and continue with the configuration
Discard your changes	Click Cancel

Configuring the RADIUS server to integrate PingID with your VPN

On the Basic Settings tab, you enable the built-in RADIUS server to integrate PingID with your VPN.

Steps

1. To configure the RADIUS client and server information, fill out the required fields.

For information about each field, see the following table.

Field	Description
Client IP	Enter the IP address of the VPN server.
Client Shared Secret	Enter the password shared between the VPN server and PingFederate.
Server Authentication	Enter the listening port of the integrated RADIUS server.
Port	The default value is 1812.

- 2. If you are connecting PingFederate to a directory server, select the Validate with LDAP check box to enable first-factor credential validation against the directory server. This check box is selected by default.
- 3. If you are connecting PingFederate to a directory server and have updated the Search Filter value to compare user input against another user attribute, enter that user attribute name in the PingID Username Attribute field. Click Next.

Example:

If you have updated the Search Filter value from uid=\${username} to mail=\${username}, enter mail in the PingID Username Attribute field.

For more information about the Search Filter field, see .

Reviewing the PingID VPN (RADIUS) configuration

On the **Summary** tab, review your configuration.

Steps

Perform one of the following actions.

Choose from:

- To edit your configuration, click the corresponding tab and follow the wizard.
- To keep your changes, click **Done** and continue with the configuration.
- To discard your changes, click Cancel.

Configuring provisioning to PingID

If you have chosen to connect PingFederate to a directory server, the Provisioning tab becomes available.

About this task

On the **Provisioning** tab, you have the option to enable provisioning of users from your directory server to PingID. In this configuration, specify the group where PingFederate should look for member users and update PingID when their email address, first name, or last name has changed.



This provisioning capability is designed to manage existing PingID accounts. It does not create new PingID users. When PingFederate detects that a user has been removed from the specified group or disabled in the directory server, PingFederate sends an update to PingID to disable the PingID service for that account.

Steps

- 1. Select the **Configure Provisioning** check box.
- 2. In the PingID Group DN field, enter the distinguished name (DN) of the applicable group.



Note:

The specified group must reside under the hierarchy of the previously-defined Search Base value. For more information, see .

3. If you want PingFederate to monitor changes for users through nested group membership, select the Nested check box. Click Next.

Entering basic information for PingFederate Bridge

On the **Basic Information** tab, enter your federation information.

Steps

1. In the **Base URL** field, verify your base URL. Update as needed.

The domain portion of the base URL should match the domain name of your organization because it is part of the address where your applications, users, and partners communicate with your PingFederate environment.



Note:

You can add multiple virtual host names at a later time. For more information, see Virtual host names on page 854.

2. In the **Entity ID** field, enter your Entity ID if prompted. Click **Next**.



Note:

This is the unique identifier of your organization. It is how your partners identify you when communicating with you based on SAML 2.0 specifications.

Reviewing the PingFederate Bridge configuration

On the **Confirmation** tab, review your configuration.

Steps

Perform one of the following actions.

Choose from:

- To amend your configuration, click the corresponding tab, and then follow the configuration wizard.
- To keep your configuration, click Next, and then click Done.
- To discard your changes, close your browser and restart PingFederate.

Completion of the PingFederate Bridge setup

The **Complete!** tab congratulates you on successfully setting up PingFederate Bridge.

Steps

- 1. Make a note of the instructions under What's Next? You will need to complete these tasks when you begin configuring PingFederate Bridge.
- 2. Click Done.

PingFederate administrative console

The PingFederate administrative console provides a wizard-like interface in which you configure your federation use cases.

The menu items depend on the administrator's permissions. For more information, see *Administrative* accounts on page 842.



Note:

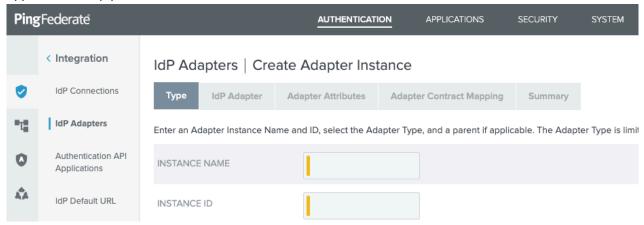
The administrative console manages the following runtime functions:

- Performing periodic configuration archive backup
- Cleaning expired persistent authentication sessions
- Cleaning expired access grants
- Updating connections from metadata URLs, including PingOne SP connections if configured, and sends email notifications
- Performing automatic rotation of signing certificates if enabled

If you prefer not to keep the administrative console running to perform these tasks, you can create your own process to manage them.

Tasks and steps

Each task consists of a series of tabs. Each tab consists of a sequence of steps. The tasks and tabs appear in the top portion of the window.



Sample tasks and steps

In this example, the primary task is managing one or more IdP adapter instances (IdP Adapters). The secondary task is creating an adapter instance (Create Adapter Instance). The current tab selects the type of adapter (Type). The subsequent tabs, which the administrator has not yet reached, are grayed out.

The administrator console displays a summary window at the end of every task, which offers the opportunity to review and make changes as needed.

Some steps provide buttons that branch to secondary tasks with multiple tabs. When the secondary tasks are complete, the administrative console returns to the primary task for the administrators to continue with the configuration.



Note:

Clicking Cancel or Done discards all unsaved changes for the tabs shown in the current task and returns you to the window from which you accessed the task.

Example

When creating a connection to a partner, the administrator might need to create a new digital signing certificate. The administrative console provides a button to begin creating a new signing certificate. When the administrator completes the task, the administrative console returns to the primary task of creating a connection to a partner.

Console buttons

The buttons at the bottom of the administrative console change depending on where you are in the configuration process.

The following table describes these buttons.

Button	Description
Save	Saves changes for all tabs in the current task and returns to thewindow from which the task or tab was accessed. This button is available only when the Save operation is valid.
Done	Marks all steps as complete for a current task, but does not save the configuration because further tasks or steps are necessary. To save your changes, click Save , or continue the configuration until you see a Save button. When creating a new service provider (SP) or identity provider (IdP) connection, click Save Draft .
Save Draft	Saves a connection's draft configuration.
Cancel	Discards all changes and returns to the window from which the current task was accessed.
Previous	Returns to the previous tab.
Next	Proceeds to the next tab if all required steps are complete in the current tab.



CAUTION:

Do not use the browser's Back, Forward, or Refresh buttons. Always use the navigation buttons in the PingFederate user interface, Previous, Next, or Done.

Third-party cryptographic solutions

You can configure PingFederate to use a hardware security module (HSM) for cryptographic material storage and operations. Standards such as the Federal Information Processing Standard (FIPS) 140-2 require the storage and processing of all keys and certificates on a certified cryptographic module.

You can also integrate PingFederate with a third-party software cryptographic solution.

Hardware security modules

When integrating with an HSM, PingFederate must be deployed with Oracle Server JRE 8. Oracle Java SE Development Kit 11 and OpenJDK 11 are not supported.

Typically, integrating with an HSM involves two steps:

- 1. Install and configure the HSM according to the manufacturer's documentation.
- 2. Follow the vendor-specific instructions to configure a new or existing PingFederate environment to use the HSM for key generation, storage, and operation.



Use HSM hybrid mode to store each relevant key and certificate on the HSM or the local trust store. This allows you to transition the storage of keys and certificates to an HSM without needing to deploy a new PingFederate environment to mirror the setup. For more information, see *Transitioning to an HSM* on page 626.



Configuring PingFederate to use an HSM for cryptographic material storage and operations might impact performance. The level of impact depends on the performance of cryptographic functionality provided by the HSM and the network latency between PingFederate and the HSM. Consult with your HSM vendor for performance tuning if you plan to use an HSM in your PingFederate deployment.

Software cryptographic solution

PingFederate supports Bouncy Castle FIPS as the provider of its Java keystore and cryptographic operations.

Supported hardware security modules

PingFederate supports multiple configurations for secure material storage and processing.

PingFederate supports the following modules:

- AWS CloudHSM (stores private keys only)
- Gemalto SafeNet Luna Network HSM (stores private keys only)
- nCipher nShield Connect HSM (stores certificates and private keys)

Integrating with AWS CloudHSM

PingFederate supports multiple hardware security modules (HSMs), including AWS CloudHSM.

Before you begin

Ensure Oracle Server JRE 8 is installed on the PingFederate server. For more information, see *Installing* Java on page 119.

Additionally, PingFederate must be deployed on one of the Linux operating systems supported by both AWS CloudHSM and PingFederate. For more information, see System requirements on page 110 and Install and Configure the AWS CloudHSM Client (Linux) in the AWS CloudHSM documentation.

Steps

1. Request a crypto user (CU) account from your AWS CloudHSM administrator.



Note:

You need this account's username and password for your PingFederate installation.

2. Install and configure your AWS CloudHSM client and tools.



Note:

Commands vary depending on the operating system. For more information, see *Install the AWS* CloudHSM Client and Command Line Tools in the AWS CloudHSM documentation.

3. Install the AWS CloudHSM software library for Java.



Note:

Commands vary depending on the Linux operating system. For instructions, see *Installing Java Library* in the AWS CloudHSM documentation.

a. Using the command line, export four environment variables.

```
$ export LD_LIBRARY_PATH=/opt/cloudhsm/lib
$ export HSM_PARTITION=PARTITION_1
$ export HSM_USER=<HSM username>
$ export HSM_PASSWORD=<password>
```

b. To validate the AWS CloudHSM client installation, run the following Java program.

```
$ java -classpath "/opt/cloudhsm/java/*" org.junit.runner.JUnitCore
TestBasicFunctionality
```

Result:

The result should be similar to the following output.

```
JUnit version 4.12
.2020-02-24 16:17:01,681 DEBUG [main] TestBasicFunctionality
 (TestBasicFunctionality.java:33) - Adding provider.
2020-02-24 16:17:01,749 DEBUG [main] TestBasicFunctionality
 (TestBasicFunctionality.java:42) - Logging in.
2020-02-24 16:17:01,749 INFO [main] cfm2.LoginManager
 (LoginManager.java:238) - Looking for credentials in
 HsmCredentials.properties
2020-02-24 16:17:01,750 INFO [main] cfm2.LoginManager
 (LoginManager.java:256) - Looking for credentials in
 System.properties
2020-02-24 16:17:01,750 INFO [main] cfm2.LoginManager
 (LoginManager.java:264) - Looking for credentials in System.env
2020-02-24 16:17:01,784 DEBUG [main] TestBasicFunctionality
 (TestBasicFunctionality.java:54) - Generating AES Key with key size
 256.
2020-02-24 16:17:01,995 DEBUG [main] TestBasicFunctionality
 (TestBasicFunctionality.java:63) - Encrypting with AES Key.
2020-02-24 16:17:02,005 DEBUG [main] TestBasicFunctionality
 (TestBasicFunctionality.java:84) - Deleting AES Key.
2020-02-24 16:17:02,006 DEBUG [main] TestBasicFunctionality
 (TestBasicFunctionality.java:92) - Logging out.
Time: 0.334
OK (1 test)
```

For more information, see *Validating the Installation* in the AWS CloudHSM documentation.

5. Update the JAVA_HOME/jre/lib/security/java.security file in your Java environment, and then add the AWSCloudHSMProvider line to the list of security providers immediately after the sun.security.provider.Sun provider.

Example:

```
# List of providers and their preference orders (see above):
security.provider.1=sun.security.provider.Sun
security.provider.2=com.pingidentity.crypto.AWSCloudHSMProvider
security.provider.3=sun.security.rsa.SunRsaSign
security.provider.4=sun.security.ec.SunEC
security.provider.5=com.sun.net.ssl.internal.ssl.Provider
security.provider.6=com.sun.crypto.provider.SunJCE
security.provider.7=sun.security.jgss.SunProvider
security.provider.8=com.sun.security.sasl.Provider
security.provider.9=org.jcp.xml.dsig.internal.dom.XMLDSigRI
security.provider.10=sun.security.smartcardio.SunPCSC
```

6. Configure a new PingFederate installation on the network interconnected to the HSM.



Go to the next step to integrate an existing PingFederate installation with your HSM.

- 7. To enable the Java interface and PingFederate integration, copy the following files to the JAVA HOME/ jre/lib/ext directory:
 - <pf install>/pingfederate/lib-ext/pf-aws-cloud-hsm-wrapper.jar
 - <pf install>/pingfederate/server/default/lib/disruptor.jar
 - All files under the /opt/cloudhsm/java directory
 - All files under the /opt/cloudhsm/lib directory
- 8. Update the hivemodule.xml file.
 - a. Edit the <pf install>/pingfederate/server/default/conf/META-INF/ hivemodule.xml file.
 - b. Go to the <!-- Crypto provider --> section.
 - c. Update the class attribute value of the construct element for both the JCEManager and CertificateService service endpoint.

```
<!-- Crypto provider -->
<service-point id="JCEManager"</pre>
interface="com.pingidentity.crypto.JCEManager">
<invoke-factory>
 <construct class="com.pingidentity.crypto.AWSCloudHSMJCEManager"/>
</invoke-factory>
</service-point>
<service-point id="CertificateService"</pre>
 interface="com.pingidentity.crypto.CertificateService">
<invoke-factory>
  . . .
 <construct
class="com.pingidentity.crypto.AWSCloudHSMCertificateServiceImpl"/>
</invoke-factory>
</service-point>
```

- 9. Update the <pf install>/pingfederate/bin/run.properties file.
 - a. Change the value of the pf.hsm.mode property from OFF to AWSCLOUDHSM.
 - b. If you are setting up a new PingFederate installation, set the value of the pf.hsm.hybrid property to false to store newly-created or imported certificates on your HSM.
 - c. If you are configuring an existing PingFederate installation, set the value to true for the flexibility to store each relevant key and certificate on the HSM or the local trust store.
 - This allows you to transition the storage of keys and certificates to your HSM without deploying a new PingFederate environment. For more information, see *Transitioning to an HSM* on page 626.
- 10. Run the <pf install>/pingfederate/bin/hsmpass.sh script for Linux.
 - a. Enter the password for the CU account when prompted. See step 1. This procedure securely stores the password for communication to the HSM from PingFederate.

11. If the username of the CU account is not crypto user, update the

com.pingidentity.crypto.AWSCloudHSM.xml file.

a. Edit the <pf install>/pingfederate/server/default/data/config-store/ com.pingidentity.crypto.AWSCloudHSM.xml file.

The unmodified version of the com.pingidentity.crypto.AWSCloudHSM.xml file follows.

```
<?xml version="1.0" encoding="UTF-8"?>
<con:config xmlns:con="http://www.sourceid.org/2004/05/config">
   <con:item name="Partition">PARTITION 1</con:item>
   <con:item name="CryptoUser">crypto user</con:item>
</con:config>
```

b. Replace crypto user with the username of the CU account.

Example:

In the following example, the username of the CU account is example_user.

```
<?xml version="1.0" encoding="UTF-8"?>
<con:config xmlns:con="http://www.sourceid.org/2004/05/config">
   <con:item name="Partition">PARTITION 1</con:item>
    <con:item name="CryptoUser">example_user</con:item>
</con:config>
```

- 12. Repeat these steps on each node.
- 13. Start the new PingFederate server or restart the existing PingFederate server.

AWS CloudHSM operational notes

When using a hardware security module (HSM), some restrictions apply to PingFederate.

- PingFederate requires Oracle Server JRE 8 for deployment.
- PingFederate does not store public certificates on the hardware module for signature verification, encryption, and back-channel authentication. Instead, the local trust store on the file system stores these certificates.
- As an OpenID Provider, PingFederate can use static or dynamically-rotating keys to sign ID tokens, JSON web tokens (JWTs) for client authentication, and OpenID Connect request objects. When using dynamically-rotating keys as part of the default configuration, the memory, not the HSM, stores shortterm keys. The HSM can store static keys.
- Private keys are not exportable. When configured for use with the HSM, PingFederate disables administrative-console options for this feature. Only the public portion of generated keys is exportable.
- When using the Configuration Archive feature, any keys, certificates, or objects generated and stored on the HSM prior to saving a configuration archive must continue to exist unaltered when the archive is restored. In other words, the PingFederate user interface must execute any deletion or creation of objects on the HSM for proper operation.

For example, you create and save objects A, B, and C to the HSM and create a data archive that contains references to those objects. If you delete object C and attempt to recover it through the data archive, PingFederate fails. Because the data archive contains a reference to the object and the object has been deleted from the HSM, you cannot use that data archive again.

 PingFederate limits cipher suites to those listed in the <pf install>/pingfederate/server/ default/data/config-store/com.pingidentity.crypto.AWSCloudHSMJCEManager.xml file.

Integrating with Gemalto SafeNet Luna Network HSM

PingFederate supports multiple hardware security modules (HSMs), including Gemalto SafeNet Luna Network HSMs.

Steps

- 1. Ensure the PingFederate server has Oracle Server JRE 8 installed.
 - To use larger key sizes, enable the Java Cryptography Extension (JCE) "unlimited strength" jurisdiction policy. For more information, see *Installing Java* on page 119.
- 2. Install and configure your Gemalto SafeNet Luna Network HSM, including the optional JSP package for Java, according to SafeNet's instructions.

This includes creating a partition, creating a Network Trust Link (NTL), and assigning a client to a partition.

- a. Ensure the operation of the vtl verify command to indicate secure and proper communication with the HSM.
- b. Delete any unnecessary keys or objects created while testing communication to the HSM from the host running PingFederate.
- c. For your PingFederate installation, record the password used to open communication to the HSM through the NTL.
- 3. To enable the Java interface, copy the Luna library and program files to the Java installation as follows.

Operating system	Steps
Windows	Copy the LUNA_HOME\jsp\lib\LunaAPI.dll file to an arbitrary directory and add the directory's path as a system variable. Alternatively, you can copy the file to the Windows system directory (C:\Windows\System32).
	Copy the LUNA_HOME\jsp\lib\LunaProvider.jar file to the JAVA_HOME\jre\lib\ext directory.
Linux	Copy the libLunaAPI.so and LunaProvider.jar files from the LUNA_HOME/jsp/lib directory to the JAVA_HOME/jre/lib/ext directory.

Prior to installing PingFederate, SafeNet provides sample Java applications to test that the Java HSM interface works. For more information, see the HSM documentation from Gemalto.

4. Update the JAVA HOME/jre/lib/security/java.security file in your Java environment and add the LunaProvider line to the bottom of the list of security providers.

Example:

```
# List of providers and their preference orders (see above):
security.provider.1=sun.security.provider.Sun
security.provider.2=sun.security.rsa.SunRsaSign
security.provider.3=sun.security.ec.SunEC
security.provider.4=com.sun.net.ssl.internal.ssl.Provider
security.provider.5=com.sun.crypto.provider.SunJCE
security.provider.6=sun.security.jgss.SunProvider
security.provider.7=com.sun.security.sasl.Provider
security.provider.8=org.jcp.xml.dsig.internal.dom.XMLDSigRI
security.provider.9=sun.security.smartcardio.SunPCSC
security.provider.10=com.safenetinc.luna.provider.LunaProvider
```

5. On the network interconnected to the HSM, set up a new PingFederate installation.



To integrate an existing PingFederate installation with your HSM, skip to the next step.

- 6. Update the hivemodule.xml file.
 - a. Edit the <pf install>/pingfederate/server/default/conf/META-INF/ hivemodule.xml file.
 - b. Go to the <!-- Crypto provider --> section.
 - c. Update the class attribute value of the construct element for both the JCEManager and CertificateService service endpoint as follows.

```
<!-- Crypto provider -->
<service-point id="JCEManager"</pre>
interface="com.pingidentity.crypto.JCEManager">
<invoke-factory>
 <construct class="com.pingidentity.crypto.LunaJCEManager"/>
</invoke-factory>
</service-point>
<service-point id="CertificateService"</pre>
 interface="com.pingidentity.crypto.CertificateService">
<invoke-factorv>
  . . .
 <construct class="com.pingidentity.crypto.LunaCertificateServiceImpl"/>
</invoke-factory>
</service-point>
```

- 7. Update the <pf install>/pingfederate/bin/run.properties file.
 - a. Change the value of pf.hsm.mode from OFF to LUNA.
 - b. To configure a new PingFederate installation, set the value of pf.hsm.hybrid to false. When set to false, the HSM stores newly created or imported certificates.
 - To configure an existing PingFederate installation, set the value to true for the flexibility to store each relevant key and certificate on the HSM or the local trust store. This allows you to transition the storage of keys and certificates to your HSM without deploying a new PingFederate environment. For more information, see *Transitioning to an HSM* on page 626.
 - c. Add the following line at the end of the file:

```
xmlsecurity.use.oaep.parameterspec=true
```

- 8. From the <pf install>/pingfederate/bin directory, run the hsmpass.bat batch file for Windows or the hsmpass.sh script for Linux.
 - a. Enter the NTL password when prompted. For more information, see step 2.

This procedure securely stores the password for NTL communication to the HSM from PingFederate.



Note:

The Gemalto SafeNet Luna Network HSM supports configuration in a high-availability group. For more information, see the SafeNet distributed-installation instructions. To properly synchronize data, ensure that the HAOnly property is enabled using the vtl haAdmin -HAOnly -enable command.

- 9. Repeat these steps on each node.
- 10. Start the new PingFederate server or restart the existing PingFederate server.



Important:

Whenever you restart the Luna HSM, SafeNet recommends you also restart dependent processes such as PingFederate and all server nodes in a cluster.

SafeNet Luna Network HSM operational notes

When using a hardware security module (HSM), some restrictions apply to PingFederate.

- PingFederate requires Oracle Server JRE 8 for deployment.
- PingFederate does not store public certificates on the hardware module for signature verification, encryption, and back-channel authentication. Instead, the local trust store on the file system stores these certificates.
- As an OpenID Provider, PingFederate can use static or dynamically-rotating keys to sign ID tokens, JSON web tokens (JWTs) for client authentication, and OpenID Connect request objects. When using dynamically-rotating keys as part of the default configuration, the memory, not the HSM, stores shortterm keys. The HSM can store static keys.
- Private keys are not exportable. When configured for use with the HSM, PingFederate disables administrative-console options for this feature. Only the public portion of generated keys is exportable.
- When running in FIPS 140-2 level 3 compliance, also called strict FIPS mode, private keys cannot be imported. In this mode, administrative-console options for this feature are disabled.
- When using the Configuration Archive feature, any keys, certificates, or objects generated and stored on the HSM prior to saving a configuration archive must continue to exist unaltered when the archive is restored. In other words, the PingFederate user interface must execute any deletion or creation of objects on the HSM for proper operation.
 - For example, you create and save objects A, B, and C to the HSM and create a data archive that contains references to those objects. If you delete object C and attempt to recover it through the data archive, PingFederate fails. Because the data archive contains a reference to the object and the object has been deleted from the HSM, you cannot use that data archive again.
- PingFederate limits cipher suites to those listed in the <pf install>/pingfederate/server/ default/data/config-store/com.pingidentity.crypto.LunaJCEManager.xml file.

Integrating with nCipher nShield Connect HSM

PingFederate supports multiple hardware security modules (HSMs), including nCipher nShield Connect HSM.

Steps

- 1. Ensure the PingFederate server has Oracle Server JRE 8 installed. To use larger key sizes, enable the Java Cryptography Extension (JCE) unlimited strength jurisdiction policy. For more information, see *Installing Java* on page 119.
- 2. Install and configure your nCipher nShield Connect HSM client software.
 - As part of the installation, install the optional Java Support (including KeySafe) and nCipherKM JCA/ JCE provider classes components.

3. After your installation, see the HSM documentation from nCipher to make your PingFederate server a client of an HSM server.



Note:

supports both Operator Card Set (OCS) protected keys and module-protected keys.

For OCS, note the password. You need the password for your installation of .

For module-protected keys, edit the pingfederate/server/default/data/config-store/ com.pingidentity.crypto.NCipherSettings.xml file to add the following entries:

```
<con:item name="protect">module</con:item>
<con:item name="ignorePassphrase">true</con:item>
```

4. To enable the Java interface, copy the NFAST HOME/java/classes/nCipherKM.jar file to the JAVA HOME/jre/lib/ext directory.

Prior to installing PingFederate, nCipher offers sample Java applications to test that the Java HSM interface works. For more information, refer to the HSM documentation from nCipher.

5. Update the JAVA HOME/jre/lib/security/java.security file in your Java environment and add the nCipherKM line to the list of security providers, after the sun providers.

```
# List of providers and their preference orders (see above):
security.provider.1=sun.security.provider.Sun
security.provider.2=sun.security.rsa.SunRsaSign
security.provider.3=sun.security.ec.SunEC
security.provider.4=com.sun.net.ssl.internal.ssl.Provider
security.provider.5=com.sun.crypto.provider.SunJCE
security.provider.6=sun.security.jgss.SunProvider
security.provider.7=com.sun.security.sasl.Provider
security.provider.8=org.jcp.xml.dsig.internal.dom.XMLDSigRI
security.provider.9=sun.security.smartcardio.SunPCSC
security.provider.10=sun.security.mscapi.SunMSCAPI
security.provider.11=com.ncipher.provider.km.nCipherKM
```

6. Set up a new PingFederate installation on the network interconnected to the HSM.



Important:

Skip to the next step to integrate an existing PingFederate installation with your HSM.

- 7. Update the hivemodule.xml file.
 - a. Edit the <pf install>/pingfederate/server/default/conf/META-INF/ hivemodule.xml file.
 - b. Look for the <!-- Crypto provider --> section.
 - c. Update the class attribute value of the construct element for both the JCEManager and CertificateService service endpoint.

```
<!-- Crypto provider -->
<service-point id="JCEManager"</pre>
 interface="com.pingidentity.crypto.JCEManager">
 <invoke-factory>
  <construct class="com.pingidentity.crypto.NcipherJCEManager"/>
 </invoke-factory>
</service-point>
```

```
<service-point id="CertificateService"</pre>
interface="com.pingidentity.crypto.CertificateService">
<invoke-factory>
  <construct class="com.pingidentity.crypto.NcipherCertificateServiceImpl"/>
</invoke-factory>
</service-point>
```

- 8. Update the <pf install>/pingfederate/bin/run.properties file.
 - a. Change the value of pf.hsm.mode from OFF to NCIPHER.
 - b. If you are configuring a new PingFederate installation, set the value of pf.hsm.hybrid to false to store newly created or imported certificates on your HSM.
 - c. If you are configuring an existing PingFederate installation, set the value to true, which provides the flexibility to store each relevant key and certificate on the HSM or the local trust store. This capability allows you to transition the storage of keys and certificates to your HSM without the need to deploy a new PingFederate environment and to mirror the setup. For more information, see Transitioning to an HSM on page 626.
- 9. From the <pf install>/pingfederate/bin directory, run the hsmpass.bat batch file for Windows or the hsmpass.sh script for Linux.

Enter the Operator Card Set password when prompted. See step 3.

This procedure securely stores the password for communication to the HSM from PingFederate.

10. If you are setting up a new or configuring an existing PingFederate cluster, repeat these steps on each node.

When finished, use the following steps to replicate nShield data to the connected nodes in the cluster.

- a. On the console node, go to the <pf install>/pingfederate/server/default/data directory and create a sub directory named ncipher-kmdata-local.
- b. Copy to the ncipher-kmdata-local directory all files from the NFAST KMDATA\local directory, where NFAST KMDATA is an environment variable created during the nShield Connect installation.

For example, NFAST KMDATA could be set to C:\ProgramData\nCipher\Key Management Data.

c. Create a new environment variable named NFAST KMLOCAL and set it to <pf install>/ pingfederate/server/default/data/ncipher-kmdata-local.



Note:

You must define this environment variable on all servers within the cluster.

- d. Restart the nShield Connect hardserver on all PingFederate servers in the cluster. For instructions on restarting the hardserver, see the HSM documentation from nCipher.
- e. Sign on to the PingFederate administrative console and go to System # Server # Cluster Management.
- f. To push the configuration changes, including the nShield data, to the engine nodes, click Replicate Configuration.
- 11. Start the new PingFederate server or restart the existing PingFederate server.



Important:

Whenever you restart the nShield HSM, restart PingFederate and all server nodes in a cluster.

nShield Connect HSM operational notes

Some restrictions apply to PingFederate when using a hardware security module (HSM).

- PingFederate requires Oracle Server JRE 8 for deployment.
- When integrating PingFederate with nCipher nShield Connect on a platform with Oracle Server JRE 8u102, runtime errors might occur when handling certificates with a signing algorithm of RSA SHA256, SHA384, or SHA512. To resolve these runtime errors, upgrade to Oracle Server JRE 8u112.
- PingFederate only supports Operator Card Set (OCS) protected keys. If you use a standard, non-persistent OCS, removing the card from the smart card reader causes the HSM to remove the protected keys from its memory. Requests will likely fail because almost all requests require cryptographic processing. To resume operations, insert the card into the smart card reader and then restart PingFederate.

Alternatively, use a persistent OCS so that protected keys remain in memory even after the card is removed from the smart card reader. PingFederate will continue to process requests and to load keys and certificates from the HSM as needed. Until the card is inserted back into the HSM, the HSM will not support new key and certificate creation and storage. However, using a persistent OCS does not require a restart of PingFederate in this situation. For more information about persistent OCS, consult your HSM vendor.

- As an OpenID Provider, PingFederate can use static or dynamically-rotating keys to sign ID tokens. JSON web tokens (JWTs) for client authentication, and OpenID Connect request objects. When using dynamically-rotating keys as part of the default configuration, the memory, not the HSM, stores shortterm keys. The HSM can store static keys.
- Private keys are not exportable. When configured for use with the HSM, PingFederate disables administrative-console options for this feature. Only the public portion of generated keys is exportable.
- When running in FIPS 140-2 level 3 compliance, also called strict FIPS mode, private keys cannot be imported. In this mode, administrative-console options for this feature are disabled.
- When using the Configuration Archive feature, any keys, certificates, or objects generated and stored on the HSM prior to saving a configuration archive must continue to exist unaltered when the archive is restored. In other words, the PingFederate user interface must execute any deletion or creation of objects on the HSM for proper operation.

For example, you create and save objects A, B, and C to the HSM and create a data archive that contains references to those objects. If you delete object C and attempt to recover it through the data archive, PingFederate fails. Because the data archive contains a reference to the object and the object has been deleted from the HSM, you cannot use that data archive again.

 PingFederate limits cipher suites to those listed in the <pf install>/pingfederate/server/ default/data/config-store/com.pingidentity.crypto.NcipherJCEManager.xml file.

Supported software security package

PingFederate supports the Bouncy Castle FIPS provider software security package.

Integrating with Bouncy Castle FIPS provider

In Bouncy Castle FIPS mode, whenever PingFederate uses FIPS-approved algorithms, it uses the Bouncy Castle implementation of those algorithms. There are still a number of cases where PingFederate uses algorithms that are not FIPS-approved. For details on the contexts where PingFederate uses algorithms that are not FIPS-approved, contact customer support.

The integration of Bouncy Castle FIPS provider supports two phases:

- Hybrid to transition private keys from default keystore to the Bouncy Castle keystore.
- Non-Hybrid to start storing private keys only in the Bouncy Castle keystore.

Several properties in the <pf install>/pingfederate/bin/run.properties file allow you to configure these phases as shown in the following table.

Phase	Properties
Hybrid	pf.hsm.mode=BCFIPS
	pf.hsm.hybrid=true
Non-Hybrid	pf.hsm.mode=BCFIPS
	pf.hsm.hybrid=false

You can run either Java 8 or 11 when integrating with the BCFIPS provider. The setup steps are the same for both environments.



Important:

The only way to switch from BCFIPS mode back to non-BCFIPS mode is to roll back PingFederate with an archive.

Setting up with Java 8

This procedure describes how to integrate PingFederate with Bouncy Castle FIPS provider if you are running Java 8.

Steps

- 1. Download the Bouncy Castle FIPS library bc-fips-1.0.2.jar file from https:// www.bouncycastle.org/fips-java/.
- 2. Place the .jar file into your JRE; for example, JAVA HOME/jre/lib/ext.
- 3. Edit the JAVA HOME/jre/lib/security/java.security file, and place the org.bouncycastle.jcajce.provider.BouncyCastleFipsProvider line at the top of the list.

```
$JAVA HOME/jre/lib/security/java.security provider list:
security.provider.1=org.bouncycastle.jcajce.provider.BouncyCastleFipsProvider
security.provider.2=sun.security.provider.Sun
security.provider.3=sun.security.rsa.SunRsaSign
security.provider.4=sun.security.ec.SunEC
security.provider.5=com.sun.net.ssl.internal.ssl.Provider
security.provider.6=com.sun.crypto.provider.SunJCE
security.provider.7=sun.security.jgss.SunProvider
security.provider.8=com.sun.security.sasl.Provider
security.provider.9=org.jcp.xml.dsig.internal.dom.XMLDSigRI
security.provider.10=sun.security.smartcardio.SunPCSC
```

- 4. Edit the <pf install>/pingfederate/server/default/conf/META-INF/hivemodule.xml
 - a. Go go the <!--Crypto provider --> section.
 - b. Update the class attribute value of the construct element for both the JCEManager and CertificateService service endpoint.

```
<!-- Crypto provider -->
<service-point id="JCEManager"</pre>
interface="com.pingidentity.crypto.JCEManager">
   <invoke-factory>
       <construct class="com.pingidentity.crypto.BCFIPSJCEManager"/>
   </invoke-factory>
</service-point>
<service-point id="CertificateService"</pre>
interface="com.pingidentity.crypto.CertificateService">
```

```
<invoke-factory model="autoreloadable">
       <construct
 class="com.pingidentity.crypto.BCFIPSCertificateServiceImpl"/>
  </invoke-factory>
</service-point>
```

- 5. Edit the <pf install>/pingfederate/bin/run.properties file.
 - a. Change the org.bouncycastle.fips.approved only property to false.



This version of PingFederate does not support the approved_only mode because in some cases PingFederate uses algorithms that are not FIPS-approved. For information about those cases, contact customer support.

- b. Change the pf.hsm.mode property to BCFIPS.
- c. If you are setting up a new PingFederate installation, set the value of the pf.hsm.hybrid property to false to store newly created or imported certificates on your HSM.
- d. If you are configuring an existing PingFederate installation, set the pf.hsm.hybrid value to true for the flexibility to store each relevant key and certificate on the HSM or the local trust

This allows you to transition the storage of keys and certificates to your HSM without deploying a new PingFederate environment. For more information, see *Transitioning to an HSM* on page 626.

6. If you run PingFederate as a service on Windows, edit the install/pingfederate/sbin/ wrapper/PingFederateService.conf file, and remove or comment out the following line:

```
wrapper.java.additional.14 = - Dorg.bouncycastle.fips.approved only = true
```

Setting up with Java 11

This procedure describes how to integrate PingFederate with Bouncy Castle FIPS provider if you are running Java 11.

Steps

- 1. Download the Bouncy Castle FIPS library bc-fips-1.0,2.jar file from https:// www.bouncycastle.org/fips-java/.
- 2. Place the .jar file into <pf install>/pingfederate/lib.
- 3. Edit the JAVA HOME/conf/security/java.security file, and place the org.bouncycastle.jcajce.provider.BouncyCastleFipsProvider line at the top of the list.

```
$JAVA HOME/conf/security/java.security provider list:
security.provider.1=org.bouncycastle.jcajce.provider.BouncyCastleFipsProvider
security.provider.2=SUN
security.provider.3=SunRsaSign
security.provider.4=SunEC
security.provider.5=SunJSSE
security.provider.6=SunJCE
security.provider.7=SunJGSS
security.provider.8=SunSASL
security.provider.9=XMLDSig
security.provider.10=SunPCSC
security.provider.11=JdkLDAP
security.provider.12=JdkSASL
security.provider.13=SunPKCS11
```

- 4. Edit the <pf install>/pingfederate/server/default/conf/META-INF/hivemodule.xml file.
 - a. Go go the <!--Crypto provider --> section.
 - b. Update the class attribute value of the construct element for both the JCEManager and CertificateService service endpoint.

```
<!-- Crypto provider -->
<service-point id="JCEManager"</pre>
interface="com.pingidentity.crypto.JCEManager">
  <invoke-factory>
       <construct class="com.pingidentity.crypto.BCFIPSJCEManager"/>
   </invoke-factory>
</service-point>
<service-point id="CertificateService"</pre>
interface="com.pingidentity.crypto.CertificateService">
   <invoke-factory model="autoreloadable">
       <construct
class="com.pingidentity.crypto.BCFIPSCertificateServiceImpl"/>
   </invoke-factory>
</service-point>
```

- 5. Edit the <pf install>/pingfederate/bin/run.properties file.
 - a. Change the org.bouncycastle.fips.approved only property to false.



Note:

This version of PingFederate does not support the approved_only mode because in some cases PingFederate uses algorithms that are not FIPS-approved. For information about those cases, contact customer support.

- b. Change the pf.hsm.mode property to BCFIPS.
- c. If you are setting up a new PingFederate installation, set the value of the pf.hsm.hybrid property to false to store newly created or imported certificates on your HSM.
- d. If you are configuring an existing PingFederate installation, set the pf.hsm.hybrid value to true for the flexibility to store each relevant key and certificate on the HSM or the local trust store.

This allows you to transition the storage of keys and certificates to your HSM without deploying a new PingFederate environment. For more information, see Transitioning to an HSM on page

e. To ensure only elliptic curve groups supported by Bouncy Castle are used in TLS, add the following lines at the end of the file:

```
jdk.tls.namedGroups=secp256r1, secp384r1, secp521r1, \
                    ffdhe2048, ffdhe3072, ffdhe4096, \
                    ffdhe6144, ffdhe8192
```

- 6. Edit the <pf install>/pingfederate/bin/run.sh or run.bat file (depending on your operating system).
 - a. Add the following line:

```
pf fips="$PF HOME/lib/bc-fips-1.0.2.jar"
```

b. Set PF BOOT CLASSPATH to "\$runjar:\$pfrunjar:\$jettystartjar: \$pf console util:\$xmlbeans:\$pfxml:\$pf crypto luna:\$pf fips".

- 7. If you run PingFederate as a service on Windows, edit the <pf install>/pingfederate/sbin/ wrapper/PingFederateService.conf file.
 - a. Add the following line to the list of wrapper.java.classpath entries.

```
wrapper.java.classpath.10=../../lib/bc-fips-1.0.2.jar
```

b. Remove or comment out the following line.

```
wrapper.java.additional.14 = - Dorg.bouncycastle.fips.approved only = true
```

Bouncy Castle operational notes

When using the Bouncy Castle FIPS provider, some restrictions apply to PingFederate.

- As an OpenID Provider, PingFederate can use static or dynamically rotating keys to sign ID tokens, JSON web tokens (JWTs) for client authentication, and OpenID Connect request objects. When using dynamically rotating keys as part of the default configuration, the memory, not the BCFIPS key stores, stores short-term keys. The HSM can store static keys.
- PingFederate limits cipher suites to those listed in the <pf install>/ pingfederate/server/default/data/config-store/ com.pingidentity.crypto.com.pingidentity.crypto.BCFIPSJCEManager file.

Server Clustering Guide

This section introduces the server clustering functionality of PingFederate.

Overview of clustering

Server clustering allows you to define a single configuration for multiple PingFederate servers and to address single sign-on (SSO) and single logout (SLO) requests as a single system.

PingFederate includes clustering features that allow a group of PingFederate servers to appear as a single system to browsers and partner federation servers. In this configuration, all client traffic normally goes through a load balancer, which routes requests to the PingFederate servers in the cluster.

Server clustering can facilitate high availability of critical services and increase both performance and overall system throughput. However, availability and performance are often at opposite ends of the deployment spectrum, requiring administrators to balance one against the other to accommodate specific deployment goals. This topic identifies some of these choices.



Note:

PingFederate provides separate failover capabilities specifically for Outbound Provisioning, which by itself does not require either load balancing or state management. For more information, see Deploying provisioning failover on page 211.

Running multiple maintenance versions

As of PingFederate 10.0, you can run multiple maintenance versions in a cluster. For example, a cluster can contain servers running 10.0.0, 10.0.1, and 10.0.2. However, a cluster cannot contain servers running multiple major or minor versions, such as 10.0.0 and 10.1.0.

When running multiple versions in a cluster, the following message appears at the top of the console: "The cluster is running more than one version of PingFederate. Visit the Cluster Management page to see the versions." The Cluster management on page 854 menu also contains a **Version** column that displays the version of PingFederate running on each server node.

Running multiple versions of PingFederate in a cluster might cause inconsistencies in runtime behavior.

Clustering with multiple maintenance versions is meant to reduce the upgrade burden by eliminating downtime. Running in mixed mode should be a temporary solution, letting you gradually update each node until all of them are running the same maintenance version.



Important:

When upgrading the servers in the cluster, upgrade the administrative console first. The maintenance release might contain UI changes that will need to be replicated throughout the cluster.

General architecture

The cluster architecture has two layers:

Cluster-protocol layer

The cluster-protocol layer allows the PingFederate servers to discover a cluster, communicate with each other, detect and relay connectivity failures, and maintain the cluster as individual servers join and leave.

Runtime state-management services

The runtime state-management services communicate session-state information required to process SSO and logout requests. PingFederate abstracts its runtime state-management services behind Java service interfaces, which enables PingFederate to use interface implementations without regard to underlying storage and sharing mechanisms. Abstracting also provides a welldefined point of extensibility in PingFederate. Depending on the chosen runtime state-management architecture, each service can share session-state information with a subset of nodes or all nodes.

Runtime state-management architectures

PingFederate supports both adaptive clustering and directed clustering. Adaptive clustering offers the benefits of scaling PingFederate horizontally with little or no configuration requirement. Directed clustering allows administrators to specify which runtime state-management service uses which architecture model.

Group-RPC oriented approach

The prepackaged state-management implementations use a remote-procedure-call (RPC) framework for reliable group communication in a variety of deployments, allowing PingFederate servers to share state information within a cluster.

Load balancing

Clustered deployments of PingFederate for single sign-on (SSO) and logout transactions typically require the use of at least one load balancer, fronting multiple PingFederate servers.

When a client accesses the load balancer's virtual IP, the balancer distributes the request to one of the PingFederate servers in the cluster. Based on the configuration of the associated runtime-state management service, the processing server contacts other PingFederate servers through remote procedure calls as it processes SSO and logout requests.

PingFederate does not automatically balance the traffic among the servers in the cluster. You must manage SSO and logout requests externally to avoid overloading individual servers in the cluster. Because each server can only handle a certain amount of traffic, see PingFederate Performance Tuning Guide on page 988 to plan ahead.

The method that the balancer uses to select the appropriate server can vary from simple to highly complex, depending on deployment requirements. For specific balancing strategies, their strengths and weaknesses, as well as the impacts on PingFederate performance, see Runtime state-management architectures on page 185.

Load balancers can incorporate SSL/TLS accelerators or work closely with them. Due to the high computational overhead of the SSL handshake, Ping Identity recommends terminating SSL/TLS on a dedicated server external to PingFederate for deployments in which performance is a concern. You can still use SSL between the proxy or balancer and PingFederate, but as a separate connection.

Server modes

In a cluster, you can configure each PingFederate instance, or node, as either an administrative console or a runtime engine. Runtime engines (also known as engine nodes) service federated-identity protocol requests, while the console server (also known as the console node) administers policy and configuration for the entire cluster through the administrative console. A cluster can contain one or more engine nodes but only one console node.



You must set the administrative console node to run outside of the load-balanced group in order to successfully process SSO requests.

Cluster protocol architecture

PingFederate's cluster-protocol services manage discovery, cluster messaging, connectivity failure detection, membership, and merging of split clusters.



Important:

Nodes in the cluster must be able to communicate with one another over both the cluster bind port and the cluster failure detection port. This communication requirement remains true regardless of the chosen cluster discovery method or runtime state-management architecture.

Cluster discovery

PingFederate supports two cluster discovery methods.

Static discovery

Static discovery is suitable for a small cluster with about five to six engine nodes. Configuration requires no external component. You must configure each node with at least one expected node in a cluster. In practice, the initial discovery list should contain all nodes known in advance in the cluster, including itself, to increase the likelihood of new members finding and joining the cluster.

Dynamic discovery

Dynamic discovery is well-suited for environments where traffic volume may spike and require additional resources during peak hours. Instead of configuring a static list of known nodes ahead of time, configure new nodes to pull cluster membership information from a centralized repository. Because safe storage and ready accessibility of the information by all nodes is crucial, PingFederate supports IAM roles for Amazon Elastic Compute Cloud (Amazon EC2), Amazon Simple Storage Service (Amazon S3), and OpenStack Swift. The dynamic discovery method requires only a onetime setup. Once configured, maintaining a static discovery list requires no coordination effort.

Regardless of the discovery method, as individual nodes join and leave the cluster, the cluster-protocol service synchronizes the new membership information across all nodes.

Failure detection

The failure detection mechanism detects network connectivity failures by establishing TCP connections with other nodes at their cluster failure detection ports and sending occasional network messages. When a node detects a failure, it propagates the condition to other nodes, sharing new membership information across the cluster.



Important:

If you deploy any networking devices, such as a firewall, between nodes, you must configure them to allow inbound TCP connections to the cluster failure detection ports, and not to terminate these connections based on their potentially low volumes of network activities.

Runtime state-management architectures

This section provides an overview of runtime state-managment architecture options available for PingFederate.

Runtime state-management services distribute session-state information in the cluster, making it possible for multiple nodes to handle SSO and logout requests as a single system. In a cluster consisting of a large number of nodes, it might be desirable for performance reasons to share session-state information with only a subset of nodes. PingFederate supports two runtime state-management architectures: adaptive clustering and directed clustering.

Adaptive clustering

Adaptive clustering automatically distributes session-state information to multiple nodes. Administrators do not have to modify individual configuration files to specify which nodes should participate in tracking user sessions.

In essence, each session receives an address from within an internally-defined range. For redundancy, multiple nodes store each session. These nodes form a replica set. Any node that receives a request and must look up or store session-state information can do so by calculating the address of the session and reaching out to the corresponding replica set.

As individual nodes join and leave the cluster, adaptive clustering redistributes session-state information to maintain the replica set throughout the cluster.

The default size of a replica set is three, which provides redundancy in case two nodes fail and ensures that a single node's slow response time doesn't delay requests. The replication.factor setting is in the <pf install>/pingfederate/server/default/conf/cluster-adaptive.conf file.

Enable adaptive clustering by setting the pf.cluster.adaptive property in the run.properties file to true. This is the default state in new installations. For upgrades, if such property is not found or is set to false, the system disables adaptive clustering and enables directed clustering instead. To enable or disable adaptive clustering, set the pf.cluster.adaptive property to true or false on each node and then restart PingFederate. The run.properties file is in the <pf install>/pingfederate/bin directory.



Important:

After making changes to the cluster-adaptive.conf and the run.properties files, you must manually repeat the changes to all nodes in the cluster. The configuration replication process does not push these files across the cluster. When you are finished, restart PingFederate to apply the changes.



Note:

Adaptive clustering does not support the SAML 2.0 single logout (SLO) profile using the SOAP binding. If you have configured one or more SAML 2.0 connections to support SLO using SOAP, you must either share all nodes or designate state servers deployment strategies in directed clustering. For more information, see *Directed clustering* on page 189.

Other advanced settings

Fine-tune each runtime state-management service implementation separately by modifying a configuration file located in the <pf install>/pingfederate/server/default/conf directory. After making changes in these files, you must apply the changes to all nodes in the cluster manually.



The adaptive clustering concept is not applicable to the Artifact-Message Persistence and Retrieval Service, which always shares messages across all nodes to fulfill its objectives. As needed, you can modify other applicable properties, such as the rpc.timeout property. For more information, see Artifact-Message Persistence and Retrieval Service on page 198.

The following tables indicate the configuration file that applies to each implementation and the applicable properties. See the indicated sections for detailed information about each implementation.

Configuration file and service implementation

Configuration file	RPC-based service implementation
cluster-account- locking.conf	Account Locking Service on page 200
cluster- artifact.conf	Artifact-Message Persistence and Retrieval Service on page 198
cluster- assertion-replay- prevention.conf	Assertion Replay Prevention Service on page 197
cluster- idp-session- registry.conf	IdP Session Registry Service on page 196
cluster-inter- request-state.conf	Inter-Request State-Management (IRSM) Service on page 195
cluster-session- revocation.conf	Back-Channel Session Revocation Service on page 199
cluster- sp-session- registry.conf	SP Session Registry Service on page 196

Property description

Property	Description
rpc.timeout	How long, in milliseconds, this node waits before timing out unresponsive RPC invocations. The default value is 500, or half a second.

Description

synchronous.retrievelndicatesites madry responses to wait for when making synchronous remote procedure calls. When set to true, this node waits for the majority of the local replica set to respond. When set to false, it waits for all recipients to respond. true is the default value.



Note:

This property is not applicable to the Account Locking Service and not found in the cluster-account-locking.conf file.

(found only in the cluster-sessionrevocation.conf file)

bulk.revoked.sris.timeode downloads a full revocation list from another node during startup or when it rejoins a cluster after being disconnected from it, for example due to a temporary network issue. This setting determines the amount of time in milliseconds PingFederate waits before aborting the download and reporting a timeout error.

The default value is 10000, which is 10 seconds.

read.local.only

Determines how PingFederate should process queries for revocation status.

(found only in the cluster-sessionrevocation.conf file) When set to true, PingFederate processes queries for revocation status locally. When set to false, the processing node pulls revocation status from other engine nodes in the cluster, subject to the rpc.timeout value. true is the default value.



Note:

When adding a session to the revocation list, the processing node always propagates the information to all engine nodes in the cluster. For more information, see Back-Channel Session Revocation Service on page 199.



When you have enabled adaptive clustering, PingFederate ignores other properties found in these configuration files—namely preferred.node.indices and preferred.node.group.id. The latter is only in the cluster-idp-session-registry.conf file.

Multi-region support

PingFederate supports multi-region server clusters in adaptive clustering architecure.

When a cluster spans multiple regions, administrators can specify region identifiers for different groups of nodes. When regions are defined, any node that receives a request and must store session-state information can do so by sending the information to replica sets in both the local and remote regions. Requests that require read-only access to session-state information are answered locally for optimal performance.

As individual nodes in different regions join and leave the cluster, adaptive clustering redistributes sessionstate information within the region where changes in the cluster membership occur. This approach strikes a balance between minimizing the volume of session-state network traffic and improving the accuracy of session-state information across regions.

Cross-region support is enabled by default when you configure region identifiers in adaptive clustering environments. PingFederate provides cross-region support for the following functions:

- User session-state information maintained by the Inter-Request State-Management (IRSM) Service on page 195, the IdP Session Registry Service on page 196, and the SP Session Registry Service on page 196
- Assertion Replay Prevention Service on page 197
- Account Locking Service on page 200
- Replication, validation, and revocation of access tokens using the reference-token data model

When cross-region support is disabled in individual areas, engine nodes only communicate session-state information to and from the local replica set. To improve the accuracy of session-state information, you can deploy a network traffic management solution to persist, or stick, user sessions so that each subsequent request from the same user is directed to the same set of nodes.



To reduce cross-region network traffic, PingFederate does not normally replicate SSO transaction states to other regions. However, if DNS sends user requests to different regions during a single SSO transaction, the transaction will fail with the error Unable to resume processing because saved state was not found for key.

To let PingFederate asynchronously replicate SSO transaction states to other regions, open the clusteradaptive.conf file and change the value of inter.group.replicate.transaction.state to

OAuth access token management

PingFederate shares reference token information with a replica set when adaptive clustering is enabled. If region identifiers are defined, PingFederate shares reference token information among multiple replica sets across regions. Like other services, you can optionally override this default behavior by changing the inter.group.replicate.reference.tokens value in the <pf install>/pingfederate/ server/default/conf/cluster-adaptive.conf file.

When you disable cross-region support for access tokens using the reference-token data model, PingFederate does not share reference token information across regions. As a result, PingFederate cannot de-reference, validate, or revoke a reference-style access tokens issued outside of its region. For this reason, we recommended switching to the self-contained token data model prior to disabling cross-region support for the reference-token data model.

Configuring multi-region support

Define region identifiers and configure cross-region settings for multi-region PingFederate server clusters.

Steps

1. To define a region identifier for a given node, update the node.group.id value in the <pf install>/pingfederate/server/default/conf/cluster-adaptive.conf file, which is a per-server configuration.

Example: For example, if you have five engine nodes in the West Coast and six engine nodes in the East Coast, you can update the node.group.id value to W for each of the West Coast nodes and E for each of the six nodes in the East Coast.

a. Restart PingFederate after making changes to the cluster-adaptive.conf file. Result:

Once defined, the identifiers for all nodes are displayed on the System # Server # Cluster Management menu.

2. To configure cross-region support for individual areas, follow the inline instructions in the clusteradaptive.conf file to update the relevant setting values.

Directed clustering

This topic provides an overview of manual configuration of PingFederate sever nodes through directed clustering.

Directed clustering allows administrators to manually specify which PingFederate nodes should participate in tracking user sessions. Most group Remote Procedure Call (RPC)-based service implementations make use of a preferred-nodes concept, which assigns each node a list of other nodes, identified by index, with which it shares session-state information.

Each service implementation is controlled separately by a configuration file located in the <pf install>/ pingfederate/server/default/conf directory. You must replicate any changes manually for each cluster node.

The following tables indicate the configuration file that applies to each implementation and the applicable properties.



The Artifact-Message Persistence and Retrieval Service uses only the rpc.timeout setting.

Configuration file and service implementation

Configuration file	RPC-based service implementation
cluster-account- locking.conf	Account Locking Service on page 200
cluster- artifact.conf	Artifact-Message Persistence and Retrieval Service on page 198
cluster- assertion-replay- prevention.conf	Assertion Replay Prevention Service on page 197
cluster- idp-session- registry.conf	IdP Session Registry Service on page 196
cluster-inter- request-state.conf	Inter-Request State-Management (IRSM) Service on page 195
cluster-session- revocation.conf	Back-Channel Session Revocation Service on page 199
cluster- sp-session- registry.conf	SP Session Registry Service on page 196

Property description

Property	Description
preferred.node.indi	chesomma-separated list of indices identifying the nodes with which this node shares session-state information for the associated service. If left blank, this node sends session-state information to all nodes in the cluster as it processes SSO and logout requests.
	The Artifact-Message Persistence and Retrieval Service and the Back- Channel Session Revocation Service do not support this parameter.
	Ignored when adaptive clustering is enabled.
	This property has no default value.
<pre>preferred.node.grou (found only in the cluster- idp-session-</pre>	pArtalphanumeric group ID for a subcluster. If specified, each subcluster must have a unique group ID. At startup, PingFederate validates that the group ID is not already registered in the cluster by another list of preferred nodes. If the validation fails, PingFederate aborts the startup process and exits.
registry.conf file)	When the group ID is specified, the session identifier contains the information about the originating subcluster. This is helpful in deployments where PingFederate has been configured to manage authentication sessions on the Authentication # Policies # Sessions window. When an engine node receives a request to query and extend a session, it can route the request to the corresponding subcluster based on the session identifier value.
	If subclusters are configured without specifying group IDs, a request to query and extend a session is processed on the subcluster that received the revocation status request, which may be different from the subcluster where the session is being tracked. As a result, the session could reach the idle timeout sooner than expected.
	Ignored when adaptive clustering is enabled.
	This property has no default value.
rpc.timeout	How long, in milliseconds, this node waits before timing out unresponsive RPC invocations.
	The default value is 500, or half a second.
synchronous.retriev	reindicates: how one may response to wait for when making synchronous remote procedure calls. When set to true, this node waits for the majority of recipients to respond. When set to false, it waits for all recipients to respond.
	The default value is true.
(found only in the	Example 19 in the amount of time (in milliseconds) PingFederate waits before aborting the download of revocation lists and reporting a timeout error. The default value is 10000, or 10 seconds.
cluster-session-revocation.conf file)	Note: A node downloads a full revocation list from another node during startup or when it rejoins a cluster after being disconnected from it.

Property	Description
read.local.only	Determines how PingFederate processes queries for revocation status.
(found only in the cluster-session-revocation.conf file)	When set to true, queries for revocation status are processed locally. When false, the processing node pulls revocation status from other engine nodes in the cluster (subject to the rpc.timeout value).
	Note: When adding a session to the revocation list, the processing node always propagates the information to all engine nodes in the cluster. See Back-Channel Session Revocation Service on page 199 for more information. The default value is true.

Preferred node indices

Configuring the preferred.node.indices property can reduce the network communications and memory footprint. However, transaction volume and distribution can affect the resulting configuration. For more information on performance tuning, see *PingFederate Performance Tuning Guide* on page 988.

Individual services within a single cluster deployment can use different preferred nodes, meaning you can set different values for the preferred.node.indices property for each service.

Using preferred nodes can translate into a variety of deployment configurations. The following sections discuss three primary strategies to consider for meeting your network requirements:

- Sharing all nodes on page 191
- Designating state servers on page 192
- Defining subclusters on page 193



Note:

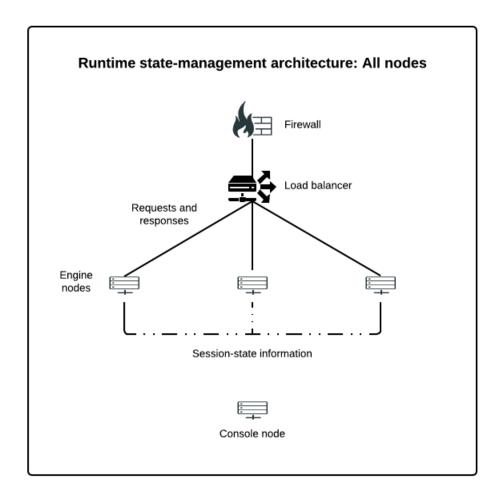
It is possible to configure overrides for authentication-adapter processing based on the runtime node servicing a request. See Configuring the Cluster Node Authentication Selector on page 223 for more information.

Sharing all nodes

The default directed clustering is the simplest deployment case, where all nodes are shared within a cluster.

Leaving the preferred.node.indices property blank in all cluster-configuration files results in a basic deployment case. An advantage of this approach is simplicity, including the option of using straightforward load-balancing strategies such as round robin. A disadvantage is that as additional nodes are added, the throughput improvement rate that clustering offers may decline as the state-replication overhead increases.

The following diagram illustrates this node-sharing approach. Requests in this deployment are directed to all nodes.



Designating state servers

A state servers clustering deployment model improves scalability by reducing communication between nodes.

You can select a few engine nodes as state servers. This deployment approach scales better than the allnodes approach because additional nodes do not require connections to every existing node, they only require a connection between each server and each state server.

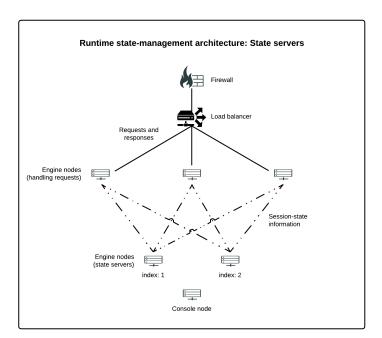
Configure this deployment by setting the preferred.node.indices of other servers in a group to those of the state servers. Configure the load balancer to isolate the state-server nodes from end-user traffic.



Note

The underlying cluster protocol still requires that all nodes are able to communicate with one another. The topology here is only an optimization for the runtime state-management services that support the concept of preferred nodes.

The following diagram illustrates the state-server approach.



In this example, the two state-server nodes have indices of 1 and 2. The preferred.node.indices property of the engine nodes handling requests would be preferred.node.indices=1,2.

Because the state servers are not processing transactions (based on the setup of the load balancer), the preferred.node.indices property for them is not used and can be left blank.



When PingFederate acts as an OAuth authorization server (AS) and the access token management instance uses a reference-token data model, the resource server (RS) must send a request to PingFederate to de-reference the access token for the corresponding identity and security information. Because the OAuth clients and the RS send their requests separately. PingFederate shares reference token information among all engine nodes despite any state server or subcluster setup.

Defining subclusters

Subclustering improves efficient scaling by limiting session-state communication to other nodes within a subcluster.

Node indices can be configured to divide a cluster into subgroups, or subclusters, of a few nodes each. Using this configuration, each node in a subcluster shares session-state information only with other members of the subcluster. This approach requires a network traffic management solution to persist, or stick, user sessions so that each subsequent request from the same user is directed to the same set of nodes. The advantage of this approach is that cluster throughput scales more linearly, because the creation of an additional subcluster will not degrade the performance of any other group.

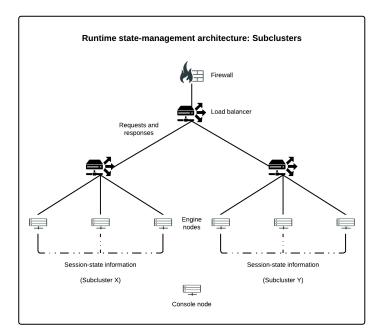


Note:

The underlying cluster protocol still requires that all nodes are able to communicate with one another. The topology here is only an optimization for the runtime state-management services that support the concept of preferred nodes.

Additionally, this architecture does not support the SAML 2.0 single logout (SLO) profile using the SOAP binding. If one or more SAML 2.0 connections are configured to support SLO via SOAP, you must choose between the sharing all nodes and designating state servers deployment strategies in directed clustering.

The following diagram illustrates the subcluster approach.



In this example, the preferred.node.indices property of each server in the cluster lists the indices of all nodes in its subgroup (including itself). Requests are directed to all nodes but the load balancer directs user sessions to the same subcluster.



Note:

When PingFederate acts as an OAuth authorization server (AS) and the access token management instance uses a reference-token data model, the resource server (RS) must send a request to PingFederate to de-reference the access token for the corresponding identity and security information. Because the OAuth clients and the RS send their requests separately, PingFederate shares reference token information among all engine nodes despite any state server or subcluster setup.

Runtime state-management services

Runtime state-management services are a collection of interfaces defining the contract that PingFederate uses to manage session states for each service.

The <pf install>/pingfederate/server/default/conf/META-INF/hivemodule.xml file specifies the implementation of each service. You only need to modify this file if you want to customize the way services are handled in a cluster. You must manually replicate your changes in each cluster node.

By default, the interfaces listed in hivemodule.xml are proxies that select the best implementation for each service, based on the operational mode of the server. For example, an in-memory-only implementation applies if the server is in standalone mode, and a group RPC-based implementation applies to servers in clustered mode. These proxies are provided for convenience. You can specifically designate your desired implementation for each service, as described in the following sections.

- Inter-Request State-Management (IRSM) Service on page 195
- IdP Session Registry Service on page 196
- SP Session Registry Service on page 196
- LRU memory management schemes on page 197
- Assertion Replay Prevention Service on page 197

- Artifact-Message Persistence and Retrieval Service on page 198
- Back-Channel Session Revocation Service on page 199
- Account Locking Service on page 200
- Other services on page 200

Configuration files for the services are located in the <pf install>/pingfederate/server/ default/conf directory.

Inter-Request State-Management (IRSM) Service

This topic is an overview of the two options for tracking session information between HTTP requests in PingFederate.

The PingFederate server tracks user-session state information between HTTP requests, such as when PingFederate, acting as an identity provider (IdP), redirects a user's browser to another system for authentication. When the user's browser returns to PingFederate after authentication, the server needs access to the state associated with that user from before the redirect. Generally, this state is short-lived.

The InterRequestStateMgmtProxy implementation chooses between two methods to track this state: group RPC-based (the clustering default) and local memory-based (the standalone default).

The configuration file is <pf install>/pingfederate/server/default/conf/cluster-interrequest-state.conf.

Group RPC-based session tracking

The group RPC-based implementation supports both adaptive clustering and directed clustering.

For adaptive clustering, PingFederate shares user session-state information with a replica set. If region identifiers are defined, PingFederate shares user session-state information among multiple replica sets across regions. You can override this default behavior in the install/pingfederate/server/ default/conf/cluster-adaptive.conf file.

For directed clustering, all preferred-node approaches are possible with this implementation.

The service-point InterRequestStateMgmt in the <pf install>/pingfederate/server/ default/conf/META-INF/hivemodule.xml file uses the proxy InterRequestStateMgmtProxy to assign this implementation as the clustering default. The specific class name is org.sourceid.saml20.service.impl.grouprpc.InterRequestStateMgmtGroupRpcImpl

Local memory-based session tracking

The local memory-based session tracking implementation tracks users in the inter-request state in the local memory of the processing server. This is the standalone default.



Important:

Adaptive clustering does not support this implementation. Use the group RPC-based session tracking instead.

The service-point InterRequestStateMgmt in the hivemodule.xml file uses the proxy InterRequestStateMgmtProxy to assign this implementation as the clustering default. The specific class name is org.sourceid.saml20.service.impl.localmemory.InterReqStateMgmtMapImpl

Local memory-based session tracking and clustering

Group RPC-based session tracking is the clustering default. To use local memorybased session tracking in a clustered environment, update the service-point InterRequestStateMgmt to use the local memory-based session tracking class, org.sourceid.saml20.service.impl.localmemory.InterRegStateMgmtMapImpl

Note:

The load balancer must support sticky sessions to force all requests for the same user session to be routed to the same server.

IdP Session Registry Service

PingFederate uses the IdP Session Registry Service to facilitate single logout (SLO) by tracking assertions issued to Service Provider (SP) partners.

PingFederate uses this service only when acting in an Identity Provider (IdP) role and supports SLO with one or more partner connections.

When PingFederate is in clustered mode, the service proxy uses a group RPC-based, preferred-nodes implementation. The configuration file is <pf install>/pingfederate/server/default/conf/ cluster-idp-session-registry.conf.

This service supports both adaptive clustering and directed clustering.

For adaptive clustering. PingFederate shares user session-state information with a replica set. If region identifiers are defined, PingFederate shares user session-state information among multiple replica sets across regions. You can optionally override this default behavior in the configuration file.

For directed clustering, all preferred-node approaches are possible with this implementation.



Both adaptive clustering and the subcluster deployment strategies in directed clustering do not support the SAML 2.0 SLO profile using the SOAP binding. If one or more SAML 2.0 connections are configured to support SLO via SOAP, you must choose between the sharing all nodes and designating state servers deployment strategies in directed clustering (see *Directed clustering* on page 189).

The service proxy uses the class

org.sourceid.saml20.service.impl.grouprpc.IdpSessionRegistryGroupRpcImpl

SP Session Registry Service

PingFederate uses the SP Session Registry Service to facilitate SLO by tracking assertions issued from identity provider (IdP) partners. PingFederate uses this service only when the server is acting in a service provider (SP) role and supports SLO with one or more partner connections.

When PingFederate is in clustered mode, the service proxy uses a group RPC-based, preferred-nodes implementation. The configuration file is <pf install>/pingfederate/server/default/conf/ cluster-sp-session-registry.conf.

This service supports both adaptive clustering and directed clustering.

For adaptive clustering, PingFederate shares user session-state information with a replica set. If region identifiers are defined, PingFederate shares user session-state information among multiple replica sets across regions. You can override this default behavior in the configuration file.

For directed clustering, all preferred-node approaches are possible with this implementation.



Note:

Both adaptive clustering and the subcluster deployment strategies in directed clustering do not support the SAML 2.0 SLO profile using the SOAP binding. If one or more SAML 2.0 connections are configured to support SLO via SOAP, you must choose between the sharing all nodes and designating state servers deployment strategies in directed clustering (see *Directed clustering* on page 189).

The service proxy uses the class

org.sourceid.saml20.service.impl.grouprpc.SpSessionRegistryGroupRpcImpl.

LRU memory management schemes

PingFederate uses an LRU memory manager to reduce memory usage by orphaned data.

During the normal course of transaction processing, the Inter-Request State-Management Service, the IdP Session Registry Service, and the SP Session Registry Service manage memory for PingFederate. End users abandoning web sessions results in orphaned data. To ensure that this orphaned data does not result in excessive memory usage, the data structures used by these services employ a least-recently-used (LRU) algorithm to purge old data. This algorithm automatically removes the oldest entries when a data structure reaches the maximum size.

Configure the maximum size of each data structure in the cpf_install/pingfederate/server/
default/conf/size-limits.conf file.

Assertion Replay Prevention Service

The Assertion Replay Prevention Service tracks POST assertions to prevent replay.

SAML standards specify that when a service provider (SP) receives assertions from the POST binding, the SP should keep track of each assertion for the duration of its validity to ensure that it is not replayed (that is, intercepted by a third party and re-posted). For OAuth and OpenID Connect, PingFederate can mandate a unique signed JSON Web Token (JWT) from the client for each request when the client is configured to authenticate via the <code>private_key_jwt</code> client authentication method, to transmit request parameters using in signed request objects, or to do both. PingFederate delegates these responsibilities to the Assertion Replay Prevention Service.

When PingFederate is in clustered mode, the service proxy uses a group RPC-based, preferred-nodes implementation. The configuration file is $< pf_install > / pingfederate / server / default / conf / cluster-assertion-replay-prevention.conf.$

The Assertion Replay Prevention Service supports both adaptive clustering and directed clustering.

For adaptive clustering, PingFederate shares token (assertion or JWT) information with a replica set. If region identifiers are defined, PingFederate shares token information among multiple replica sets across regions. You can optionally override this default behavior in the configuration file for adaptive clustering.

For directed clustering, you must choose between the sharing all nodes and designating state servers deployment strategies in directed clustering for this service.

The service proxy uses the class

org.sourceid.saml20.service.impl.grouprpc.AssertionReplayPreventionServiceGroupRpcImpl.

Unlike other services, the Assertion Replay Prevention Service fulfills only a security condition, rather than supporting normal SSO functionality, because there might be situations where the priority placed on cluster performance outweighs the priority placed on this security check. If you are in this situation, you have the option to change the implementation for the service point AssertionReplayPreventionService in the cpf_install>/pingfederate/server/default/conf/META-INF/hivemodule.xml file to one of these classes:

org.sourceid.saml20.service.impl.localmemory.AssertionReplayPreventionSvcInMemoryImpl

This is the implementation used in standalone mode. It performs all the appropriate replay checks but does not share any data with other nodes. A replay attempt routed to the same server node would fail, but other nodes would not have sufficient information to stop the transaction.

org.sourceid.saml20.service.impl.localmemory.AssertionReplayPreventionServiceNullImpl

This implementation disables assertion-replay prevention. Use with caution when performance is an absolute priority.

Artifact-Message Persistence and Retrieval Service

PingFederate's Artifact-Message Persistence and Retrieval Service keeps track of one-time keys and associated data compliant with SAML and OAuth 2.0 standards.

The following standards require PingFederate to relay data to partners using a reference-style data transportation model and to guarantee that the reference keys are valid for one-time use only.

SAML artifact binding

PingFederate sends an artifact to the partner when transmitting SAML-outbound messages using the artifact binding. Later, the partner returns to PingFederate to exchange the artifact for the actual message. If the request is valid, PingFederate delivers the message and invalidates the artifact.

OAuth 2.0 authorization grant type

When processing an authorization request from an OAuth client that uses the authorization code grant type, PingFederate returns a code to the client based on specification. The client then includes that code in its token request to PingFederate to obtain an access token. If the request is valid, PingFederate delivers the access token and invalidates the code.

The Reference ID Adapter from the Agentless Integration Kit also applies the same data transportation model and one-time-use restriction in its drop-off and pick-up operations.

In a standard environment, the PingFederate server saves the data in memory, generates a key for the data, and sends the key to the partner. The Artifact-Message Persistence and Retrieval Service keeps track of the key and the associated data until the partner contacts the PingFederate server to exchange the key for the data.

Group RPC-based retrieval

When multiple PingFederate servers are deployed to form a cluster, the keys and their data are saved in the server that creates them. Because they are not replicated to other PingFederate servers, it is possible for a key resolution request to arrive at a server that does not contain the requested data. To handle this scenario, the Artifact-Message Persistence and Retrieval Service uses a group Remote Procedure Call (RPC) retrieval approach, where the server handling the key resolution request determines the datahosting server based on the key value and contacts the appropriate server to retrieve the requested data. This group RPC implementation is compatible with the SAML artifact binding, the OAuth 2.0 authorization code grant type, and the Reference ID Adapter.



Note:

The Artifact-Message Persistence and Retrieval Service also supports a local memory approach for SAML 2.0. This approach is only suitable in clustered environments where the OAuth 2.0 authorization code grant type and the Reference ID Adapter are not in use.

When PingFederate is in clustered mode, the service proxy selects a group RPC-based implementation, which takes advantage of node indexing but not the preferred-nodes concept. Sticky-session loadbalancing strategies are not effective when the key request and its subsequent key resolution request can come from different locations.

Although this implementation does not take advantage of adaptive clustering or the preferred-nodes concept, you can configure the RPC time-out in the <pf install>/pingfederate/server/ default/conf/cluster-artifact.conf file.

SAML 2.0 indexing (local memory)

A SAML 2.0 federation entitycan support multiple artifact resolution services, each identified by a unique index number. Artifacts include this index, and a federation partner must send the artifact resolution request to the appropriate endpoint for that index. This means that servers do not need to share information concerning the artifact.

With this approach, partners must know about each of your backend servers. Generally, this means providing partners with a list that includes multiple artifact-resolution service endpoints with the corresponding indices.



PingFederate does not automatically generate this information; an administrator must create it and send it to partners who are using the artifact binding.

For example, if you have four servers in a cluster, the list might look like this:

```
<ArtifactResolutionService Binding="..." Location="https://node1/idp/</pre>
ARS.ssaml2" index="1"/>
<ArtifactResolutionService Binding="..." Location="https://node2/idp/</pre>
ARS.ssaml2" index="2"/>
<ArtifactResolutionService Binding="..." Location="https://node3/idp/</pre>
ARS.ssaml2" index="3"/>
<ArtifactResolutionService Binding="..." Location="https://node4/idp/</pre>
ARS.ssaml2" index="4"/>
```

In this case, the index corresponds to the node index configured in the run.properties file on each individual server. This service encodes the node index in the artifact handle when running in a clustered mode (it will always use an index of 0 in standalone mode).

Partners also need direct access to each ARS endpoint, which can complicate your configuration of load balancers, proxies, and firewalls. This approach cannot be used for SAML 1.x, or with adapters that utilize PingFederate's artifact-data management.

To use this approach for SAML 2.0 federation deployments, edit the <pf install>/ pingfederate/server/default/conf/META-INF/hivemodule.xml file and change the implementation for the ArtifactStore service point to the class name org.sourceid.saml20.service.impl.localmemory.ArtifactPersistenceServiceMapImpl.

Back-Channel Session Revocation Service

PingFederate uses the Back-Channel Session Revocation Service to provide OAuth clients the capabilities to add sessions to the revocation list and to guery the revocation status.

When PingFederate is in clustered mode, the service proxy uses a group remote procedure call (RPC)based implementation. When adding a session to its revocation list, the processing node always propagates the information to all engine nodes in the cluster. This allows you to choose whether queries are processed locally or after collecting information from other engine nodes.

Processing queries locally results in faster response times for engine nodes in well-connected networks. Requiring data from other engine nodes adds a layer of protection against inconsistency among engine nodes revocation lists due to network outages.

You can configure the RPC time-out and other settings in the <pf install>/pingfederate/server/ default/conf/cluster-session-revocation.conf file.

The service proxy uses the class

org.sourceid.saml20.service.impl.grouprpc.SessionRevocationServiceGroupRpcImpl.

FIFO memory management scheme

To ensure the revocation list does not result in excessive memory usage, the Back-Channel Session Revocation Service employs a first-in-first-out (FIFO) algorithm to purge old data. When the maximum size is reached, the oldest entries are automatically removed.

The maximum number of sessions is configurable by the

SessionRevocationServiceMapImpl.max.revoked.sris setting in the <pf_install>/ pingfederate/server/default/conf/size-limits.conf file. The default value is 50000.

The FIFO memory manager operates in addition to the **Session Revocation Lifetime** setting, globally configured in the **System # OAuth Settings # Authorization Server Settings** menu.

Account Locking Service

The PingFederate Account Locking Service includes account lockout prevention and password spraying prevention.

Account lockout protection prevents user accounts from locking at the underlying user repository based on too many failed authentication attempts. It also adds a layer of protection against brute force and dictionary attacks because the user is locked out for a time period when the number of failed attempts exceeds the threshold. This protection is enabled in many areas of PingFederate, including the HTML Form Adapter, the Username Token Processor, the OAuth resource owner password credentials grant type, and the native authentication scheme for the administrative console and API.

Password spraying prevention adds a layer of defense against the attack pattern where bad actors try to gain access to protected resources by using the same password, typically weak or compromised, against multiple accounts from multiple locations. When enabled, PingFederate tracks the number of failed login attempts per password. When the number of failures for a particular password reaches a threshold, that password is temporarily locked out. Password spraying prevention applies to the HTML Form Adapter, the Username Token Processor, and the OAuth 2.0 resource owner password credentials grant type.

When PingFederate is in clustered mode, the service proxy uses a group remote procedure call (RPC)-based implementation. The configuration file is $< pf_install > / pingfederate / server / default / conf/cluster-account-locking.conf.$

This service supports both the adaptive clustering and directed clustering.

For adaptive clustering, PingFederate shares state information with a replica set. If region identifiers are defined, PingFederate shares state information among multiple replica sets across regions. You can override this default behavior in the $<pf_install>/pingfederate/server/default/conf/cluster-adaptive.conf file.$

For directed clustering, PingFederate shares state information across all nodes, which helps in scenarios where PingFederate is deployed behind a load balancing infrastructure without sticky sessions.

Other services

PingFederate offers an account linking service and a pseudonym service to support SAML 2.0 federation deployment needs.

Account linking service

The account linking service stores the association between the external and internal identifiers of an end user when your implementation uses account linking as a service provider (SP) identity-mapping strategy. The default, standalone implementation uses a Java Database Connectivity (JDBC) interface to an embedded database within PingFederate. No information from the embedded database is shared across the cluster. When an identity provider (IdP) connection deployed in a cluster uses account linking, the default implementation will not work properly. In such cases, you must adjust the pointer for cluster use by pointing the service to an external database. For more information, see *Define an account-linking data store*.

Pseudonym service

The pseudonym service references the method needed by PingFederate to generate or look up a pseudonym for a user. PingFederate uses this service only if your site is acting in an IdP role and produces assertions containing pseudonyms as subject identifiers. The default implementation uses a message digest to produce the value so that no session-state synchronization is required. Developers who want to implement pseudonym handling differently can refer to the Javadoc reference describing PseudonymService interface for more information.

Deploying cluster servers

Follow these steps to configure and deploy clustered PingFederate servers.

About this task



Note:

Additional steps are required to set up failover for provisioning. If you are grouping servers exclusively to provide for provisioning failover, skip these steps and see *Deploy provisioning failover*.

Steps

- 1. Install PingFederate on each node in a cluster.
- 2. Edit the clustering properties of each node in the <pf install>/pingfederate/bin/ run.properties file. See the following table for information about each property.

Property

Description

pf.operational.modeControls the operational mode of the PingFederate server. PingFederate supports the following modes:

STANDALONE (default)

This server is a standalone instance that runs both the administrative console and runtime engine.



Important:

The value STANDALONE should only be used in a cluster where sessionstate management is not needed for any reason and configuration-archive deployment is used as the configuration synchronization method.

CLUSTERED CONSOLE

This server is part of a cluster and runs only the administration console.



Important:

Only one node in a cluster can run the administrative console.

CLUSTERED ENGINE

This server is part of a cluster and runs only the runtime engine.

Description

pf.cluster.node.ind@efines a unique index number for the server in a cluster. The index number is used to identify peers and optimize inter-node communication. The allowed range is 0 to 65535.

> If no value is set for the node index, the system assigns an auto-generated value in the range of 0 to 2147483647.

> This property has no default value. If you specify an index number, you can configure instances of the Cluster Node Authentication Selector and place them in authentication policies to customize authentication requirements based on the runtime node servicing a request.

pf.cluster.auth.pwdSets the password that each node in the cluster must use to authenticate when joining the cluster. This prevents unauthorized nodes from joining a cluster. The value can be any string, or blank.



Note:

Consider using a randomly-generated key with 22 or more alphanumeric characters. We recommend that you obfuscate the password. For information about the obfuscate command-line utility, see its built-in help.

All nodes in a cluster must share the same value, blank or otherwise.

pf.cluster.encrypt Indicates whether to encrypt network traffic sent between nodes in a cluster. The possible values are true or false (default).

> When set to true, communication within the cluster is encrypted with a symmetric key derived from the value of the pf.cluster.auth.pwd property.



Important:

When the pf.cluster.encrypt property is set to true, you must provide a value for the pf.cluster.auth.pwd property. Otherwise PingFederate aborts during its startup process.

All nodes in a cluster must have the same value for this property.

pf.cluster.encryptione kergeti sethe key that PingFederate takes into consideration when deriving the symmetric key from the value of the pf.cluster.auth.pwd property for the purpose of encrypting network traffic sent between nodes in a cluster. Required only when the pf.cluster.encrypt is set to true.

All nodes in a cluster must have the same value set for this property.

The default value is 128.

Description

pf.cluster.bind.addDefaults to NON LOOPBACK, which leaves the system to choose an available non-loopback IP address. Alternatively, enter an IP address of the network interface to which the cluster communication should bind. For machines with more than one network interface, provide a specific IP address.

> You can use this property to increase performance (particularly with UDP) and improve security by segmenting cluster-communication traffic onto a private network or VLAN.



Besides NON LOOPBACK or an IP address, you can also use other values supported by JGroups. For more information, see the bind addr special values in JGroups documentation.



Important:

This field does not support DNS name. Use the default value NON_LOOPBACK or replace it with an IP address.

pf.cluster.bind.por&pecifies the port associated with the pf.cluster.bind.address property or with the default network interface used.

> This is the port used by other cluster members during their discovery process, usually via the pf.cluster.tcp.discovery.initial.hosts property.

The default value is 7600.

pf.cluster.failure.landicentesithe.landopoptcofta server socket that is opened on the given node and used by other nodes as part of the cluster's failure-detection mechanisms. If set to 0 or unspecified, a random available port is used. The default value is 7700.

Description

pf.cluster.transportedipates the Itransport protocol used for cluster communication. Values are udp or tcp. The default value is tcp. All nodes in a cluster must have the same value set for this property.

> Use UDP when IP multicasting is enabled in the network environment and the majority of cluster traffic is point-to-full-group. You must also configure both the pf.cluster.mcast.group.address and pf.cluster.mcast.group.port properties.

Use TCP for geographically dispersed servers or when multicast is not available or disabled for some other reason. For example, when using routers that do not support multicast messaging. TCP may also be appropriate if your cluster configuration employs more point-to-point or point-to-few messaging than point-to-group. You must also configure the pf.cluster.tcp.discovery.inital.hosts property.



Note:

This property is a reference to a protocol-stack XML configuration file located in the <pf install>/pingfederate/server/default/ conf/ directory. Two stacks are provided: one for UDP multicast and one for TCP. You can customize either stack or add to it as needed by modifying the associated configuration file.

pf.cluster.mcast.grDefinacithedscaddress shared among nodes in the same cluster for UDP multicast communication; required when UDP is set as the transport protocol. The valid range is 224.0.0.0 to 239.255.255.255. Some addresses in this range are reserved for other purposes. This property is not used for TCP.

All nodes in a cluster must have the same value set for this property.

The default value is 239.16.96.69.

pf.cluster.mcast.grDafinesolHte port in conjunction with the

pf.cluster.mcast.group.address property value. This property is not used for TCP configurations.

All nodes in a cluster must have the same value set for this property.

The default value is 7601.

Property	Description	
pf.cluster.tcp.dis	for cluster membership information when discovering, joining, and rejoining the cluster. This value is required when TCP is set as the transport protocol. The value is a comma-separated list of host names (or IP addresses) and their cluster bind ports, for example, host1 [7600], 10.0.1.4 [7600], host7 [1033], 10.0.9.45 [2231].	
	When using static discovery, add at least one node for the cluster to know in advance. This property should contain all nodes in the cluster (including itself) to increase the likelihood of new members finding and joining the cluster.	
	When using dynamic discovery, leave this property blank and enable dynamic discovery in the $<\!pf_install>/pingfederate/server/default/conf/tcp.xml file. For more information, see Enabling dynamic discovery for clustering on page 206.$	
pf.cluster.adaptiv	reIndicates whether runtime state-management services should use the adaptive clustering architecture.	
	The default value is true for new installations and false for upgrades.	
pf.cluster.diagnostfæsseங்கும்sவர் JGroups diagnostics. true turns it on.		
	The default value is false.	
and	tThe multimast address and port this node listens on for diagnostic messages.	
pf.cluster.diagnos	The default values are 224.0.75.75 and 7500, respectively. Do not change the default values.	
node.tags	Defines the tags associated with this node.	
	Configuration is optional. When configured, PingFederate considers this property when processing requests. For example, you can use tags to determine the datastore location that this PingFederate node communicates with. You can also use tags in conjunction with authentication selectors and policies to define authentication requirements.	
	You can specify one tag.	
	node.tags=north	
	You can also specify a list of prioritized, space-separated tags.	
	node.tags=1 123 234	
	Tags cannot contain spaces.	

- 3. Optional: Edit configuration files in each node that control the cluster protocol and runtime state-management service. For more information, see *Runtime state-management architectures* on page 185 and *Runtime state-management services* on page 194.
- 4. Optional: If outbound provisioning is configured for your site and you want to provide failover capabilities, identify and configure the provisioning failover nodes. For more information, see *Deploy provisioning failover*.
- 5. Start or restart PingFederate on all nodes.
- 6. Sign on to the administrative console.

- 7. If you have not done so, import your PingFederate license. For more information, see *License* management on page 847.
- 8. On the System # Server # Cluster Management window, click Replicate Configuration to push the license information from the console node to all engine nodes.

Result

After the clustered environment is set up, you can start configuring PingFederate through the administrative console. When PingFederate detects a change, it prompts you to replicate the configuration to all engine nodes.

Enabling dynamic discovery for clustering

Select and configure a dynamic discovery setup for clustered PingFederate environments.

About this task

PingFederate provides five dynamic discovery choices: AWS PING, DNS PING, S3 PING, NATIVE S3 PING, and SWIFT PING.



Note:

The S3 PING discovery method has been deprecated due to AWS deprecation of the SigV2 signing method. When deployed in AWS, the suggested discovery method is NATIVE S3 PING. See the JGroups documentation for alternatives when deployed in other environments.

AWS PING enables you to scale your PingFederate infrastructure using Amazon Elastic Compute Cloud (Amazon EC2) instances in the Amazon Web Service (AWS) cloud, in one or multiple regions. PingFederate queries AWS for a list of eligible EC2 instances. If PingFederate receives at least one node, a cluster exists, and it joins that cluster. If PingFederate receives no node, it forms a new cluster. Permissions to ec2: Describe* actions must either be enabled in the AWS Identity and Access Management (IAM) role assigned to the EC2 instance or be associated with the access key parameter that you provide as part of the dynamic discovery configuration. Furthermore, you may also use a combination of tags and filters, in which case only EC2 instances that satisfy both criteria are returned.

S3 PING, SWIFT PING, and NATIVE S3 PING enable the flexibility to use both public and private cloud storage. PingFederate maintains cluster membership information in a centralized repository, a bucket in Amazon Simple Storage Service (Amazon S3) or a container in an OpenStack infrastructure. PingFederate contacts the repository for a list of nodes. If PingFederate receives at least one node, a cluster exists, and it joins the cluster and updates the repository with its information, including its IP address. If PingFederate receives no node, it forms a new cluster and updates the repository with its information so that the next node can find the new cluster. When PingFederate shuts down, it removes itself from the list and pushes an update to the repository.

NATIVE S3 PING uses AWS SDK and provides a more stable connection by using built-in security features such as obtaining credentials through IAM server instance profiles. This protocol is the recommended dynamic discovery mechanism when you are running in AWS but not using Kubernetes.

DNS PING uses DNS A or SRV records to perform discovery. This protocol is the recommended dynamic discovery mechanism when using Kubernetes. For more information, see http://www.jgroups.org/manual4/ index.html# dns ping.

You can configure dynamic discovery in the <pf install>/pingfederate/server/default/ conf/tcp.xml file. You do not need to configure the pf.cluster.tcp.discovery.initial.hosts property in the run.properties file.



Important:

You must manually configure or synchronize the dynamic discovery properties in the top.xml file on each node. The tcp.xml file is not synchronized automatically across the nodes in a cluster, nor is it part of the Replicate Configuration process. Restart PingFederate servers if they are running.

Steps

1. To configure cluster protocol properties, edit the <pf install>/pingfederate/bin/ run.properties file. Refer to the inline comments and the following table.

Property	Description
pf.operational.mode	Configure the operational mode of PingFederate. A a value of CLUSTERED_CONSOLE denotes a PingFederate administrative console node. CLUSTERED_ENGINE denotes a PingFederate runtime engine node.
	A PingFederate cluster has only one administrative console node. As you scale your PingFederate infrastructure, set the pf.operational.mode property value to CLUSTERED_ENGINE to deploy additional PingFederate runtime engine nodes.
pf.cluster.tcp.discovery.initiaRelmosesny configured value. No IP addresses are required here.	
pf.cluster.transport.protocol	Set the value to tcp.
pf.cluster.*	Refer to the inline comments to configure the rest of the pf.cluster.* property values.



Important:

You must manually configure the clustering properties on each node. The run.properties file is not copied from the console node to the engine nodes automatically, nor is it part of the Replicate Configuration process. For more information, see *Deploying cluster servers* on page 201.

2. Configure dynamic discovery properties.

Edit the <pf install>/pingfederate/server/default/conf/tcp.xml file. Refer to the inline comments and one of the following tables.

AWS_PING	
Property	Description
port_number	The port on which PingFederate listens for cluster communication.
	The default value is \${pf.cluster.bind.port}, which pulls the value defined by the pf.cluster.bind.port property in the <pf_install>/pingfederate/bin/run.properties file.</pf_install>

AWS_PING	
Property	Description
port_range	The number of additional ports that PingFederate can probe when attempting to connect to other nodes in the event that it fails to connect to the port specified by the port_number property.
	For example if the port_number property is 7600, a port_range property value of 0 means PingFederate will only try to connect at port 7600. If the port_range property value is set to 2, PingFederate will make up to two additional attempts, with the port value increasing by one each time. For example, if PingFederate fails to connect at port 7600, it will try at port 7601. If it fails again, it will try at port 7602.
	The default value is 0.
regions	A comma separated list of EC2 regions in which PingFederate will attempt discovery.
	If no regions are specified, only nodes in the same region as this node can be discovered. List all regions where you have nodes running.
	For information about regions, see the AWS documentation on <i>Regions, Availability Zones, and Local Zones</i> .
tags	A comma separated list of EC2 tag names.
	When specified, only EC2 instances that have been assigned with the specified tags can be discovered. If multiple tags are specified, only EC2 instances that have been assigned with all tags can be discovered.
	For information about tags, see the <i>documentation</i> from Amazon (docs.aws.amazon.com/AWSEC2/latest/UserGuide/Using_Tags.html).
filters	A semi-colon separated list of key value pairs of system metadata.
	When specified, only EC2 instances that match the specified filters can be discovered. If multiple filters are specified, only EC2 instances that match all filters can be discovered.
	Note that you can enter a comma separated list of values for each filter. In this case, a filter is considered a match if one of the values is satisfied. For example, if you enter: filters="instance-type=t2.small,t2.medium; architecture=x86_64", then only EC2 instances that have an instance-type of either t2.small or t2.medium and also an architecture of x86_64 can be discovered.
	For information about filters, see the <i>describe-instances documentation</i> from Amazon.
access_key an secret_key	The access key and its secret key for the purpose of querying AWS for EC2 instances.
	Applicable and required only if permissions to ec2: Describe* actions are associated with the access key.
	Keys can be encrypted using the obfuscate utility (obfuscate.bat for Windows or obfuscate.sh for Linux), located in the <pre>pf_install>/ pingfederate/bin directory.</pre>

AWS_PING	
Property	Description
log_aws_error_metasageset to true (the default), error messages received from AWS are logged to the server log.	

S3_PING	
Property	Description
location	The name of the bucket in your Amazon S3 environment.
	For information about buckets, see the <i>Create a bucket documentation</i> from Amazon.
	The security credentials to access your Amazon S3 environment.
secret_access_	For information about both keys, see the <i>Understanding and getting your</i> security credentials documentation from Amazon.
	Keys can be encrypted using the obfuscate utility (obfuscate.bat for Windows or obfuscate.sh for Linux), located in the $<\!pf_install>\!/$ pingfederate/bin directory.
remove_all_dataWhenvsetwodhang(she default), JGroups cleans up data when it detects a view change.	
	For more information, see <code>jgroups.org/manual4/index.html#FILE_PING</code> .
write_data_on_fIhe default value is true, which resolves the issue where subclusters could fail to merge after a network partition.	

SWIFT_PING	
Property	Description
auth_type	The authentication type.
auth_url	The authentication URL.
username and	The security credentials.
password	Password can be encrypted using the obfuscate utility (obfuscate.bat for Windows or obfuscate.sh for Linux), located in the $<\!pf_install>\!/$ pingfederate/bin directory.
tenant	The name of your OpenStack Keystone tenant.
container	The name of the root container.
	For more information about each of the SWIFT_PING properties, see jgroups.org/manual4/index.html#_swift_ping.

SWIFT_PING		
Property	Description	
remove_all_data_Whrenvsetwto_dhang(she default), JGroups cleans up data when it detects a vie change.		
	For more information, see <code>jgroups.org/manual4/index.html#FILE_PING</code> .	

NATIVE_S3_PING			
Property	Description		
region_name	The name of the region in which discovery should be attempted.		
	If no region is specified, only nodes in the same region as this node can be discovered.		
bucket_name	The name of the bucket in your Amazon S3 environment.		
	For information about buckets, see the <i>Create a bucket documentation</i> from Amazon.		
remove_all_data_on_view_chaWhhen set to true (the default), JGroups cleans up data when it detects a view change.			
	For more information, see http://www.jgroups.org/manual4/index.html#FILE_PING .		
write_data_on_find	Default value of true resolves the issue where subclusters could fail to merge after a network partition.		

DNS_PING		
Property	Description	
dns_query	A DNS query to the DNS Server and obtains information about the cluster members. For example, jgroups-dns-ping.myproject.svc.cluster.local.	
	For more information, see http://www.jgroups.org/manual4/index.html#_dns_ping .	

- 3. Start or restart PingFederate.
- 4. Repeat these steps for each node.



Note

You must repeat these steps for new nodes when they are added.

Next steps

After the initial setup, your nodes are ready to be deployed, undeployed, and redeployed as traffic volume changes.

Discovery mechanisms are separate from runtime state-management architectures. Discovery mechanisms determine how to find nodes to retrieve cluster information for the purpose of joining and rejoining a cluster. Runtime state-management architectures determine which nodes session-state information is shared to and fetched from.

PingFederate supports adaptive clustering and directed clustering runtime state-management architectures. When opting for dynamic discovery, consider enabling adaptive clustering whenever

possible. If multiple regions are involved, configure multi-region support for adaptive clustering as well. For more information and configuration steps, see *Adaptive clustering* on page 185.

Regardless of the chosen runtime state-management architecture, all nodes must still be able to communicate with other nodes for clustering-protocol messages. For more information, see Runtime statemanagement architectures on page 185.

Deploying provisioning failover

After configuring outbound provisioning, you can set up one or more PingFederate failover servers specifically for provisioning backup.

About this task

Provisioning runtime processing and failover is independent of single sign-on (SSO) or single logout (SLO) runtime processing and server clustering. However, if you are already deploying, or have deployed, a cluster for federation-protocol runtime processing, you can use a subset of those servers for provisioning failover. Alternatively, you can mix the configuration or set up provisioning-failover servers independently.



Note:

Each server in the failover network must be configured to use the same relational database.



CAUTION:

Use the built-in HSQLDB only for trial or training environments. For testing and production environments, always use a secured external storage solution for proper functioning in a clustered environment.

Testing involving HSQLDB is not a valid test. In both testing and production, it might cause various problems due to its limitations and HSQLDB involved cases are not supported by Pingldentity.

Steps

1. Select two or more runtime instances of PingFederate to configure for provisioning failover.

2. For each server instance, edit provisioning properties in the <pf install>/pingfederate/bin/ run.properties file as follows:

Property

Description

pf.provisioner.modeThe status of outbound provisioning. Allowed values are:

OFF (default)

Outbound provisioning is disabled.

STANDALONE

Provisioning is enabled, without failover.

FAILOVER

Provisioning is enabled, with failover.



Important:

The value STANDALONE cannot be used for failover configuration. This property must be set to FAILOVER on the primary and secondary servers.

provisioner.node.idThe unique index number of the provisioning server.

Each server must have a unique index number, which is used to prioritize which server is currently active and which is next in line in case of a failure. Values are any number.



Important:

The primary active primary server should have an index number of 1. The lowest value in the environment becomes the primary.

These node IDs are not required to start at 1, but it is recommended that they start at 1. The node IDs for each node must increase sequentially. The number must not exceed the maximum integer value supported by Java, which is 2147483647. Initial start-up performance degradation may result if the node ID does not start at 1.

grace.period

provisioner.failoveThe time interval (in seconds) between the first indication that a node is dead and failover to the next server in line. The time period should be greater than the Synchronization Frequency set in the System # Server # Protocol Settings # Outbound Provisioning tab on the administrative console.

The default value is 600, which is 10 minutes.



Important:

You must seperately configure the failover properties in the run.properties file on each provisioning server, because the run.properties file is not copied among the provisioning servers automatically or as part of the Replicate Configuration process.

- 3. Start or restart all of the PingFederate servers.
- 4. If you have not already done so, set up an external database to facilitate provisioning and then update the Internal Provisioning Data Store setting on the System # Server # Protocol Settings #

Outbound Provisioning tab. See Configuring outbound provisioning settings on page 839 for more information.

5. After configuration, if the provisioning servers belong to the same PingFederate clustered environment, go to the System # Server. In the Cluster Managementwindow, replicate the new Internal Provisioning Data Store setting to all nodes. If the provisioning servers are individual PingFederate servers, for each provisioning server, create a datastore connection to the same external database and update the Internal Provisioning Data Store setting manually.

Configuration synchronization

All nodes in a PingFederate clustered environment must have the same configuration settings, as set through the administrative console. You can use any of the following methods to ensure that configuration data is synchronized on all cluster nodes.

- Push from the administrative console.
- Deploy configuration archive.
- Make a RESTful API call to the /cluster administrative API endpoint.
- Make a web service call to the /pf-mgmt-ws/ws/ConfigReplication Connection Management Service endpoint.



Note:

Changes made directly to configuration files must be replicated manually across the cluster, as applicable. If the PingFederate servers are running, you must restart them after you replicate the changes.

Console configuration push

When multiple PingFederate servers are set up to run as a cluster, the administrative console provides a Cluster Management window.

Whenever applicable changes are made through the administrative console, a message appears at top of the console as a reminder to go to the Cluster Management window and to replicate the current console configuration to all server nodes in the cluster.



Note:

You must also use the Replicate Configuration window to initiate the transmission of the license file from the console node to all server nodes.

The Cluster Management window is also useful for verifying the current member servers of your PingFederate cluster.

Configuration-archive deployment

Uploading configuration archives is an alternate method of copying configurations to clustered PingFederate servers.

After you configure or reconfigure the console, you can also update cluster nodes by downloading a configuration archive from the System # Server # Configuration Archive window and then deploying it either manually or using a scripted process to the <pf install>/pingfederate/server/default/ data/drop-in-deployer directory on each cluster node or provisioning-failover server.

A configuration archive contains the same information sent during the configuration push from the administrative console described in Console configuration push on page 213. However, configurationarchive deployment also provides for scheduling and scripting cluster synchronization.

If you have configured one of the following runtime state-management services on the engine nodes, you must manually migrate the configuration files to the engine nodes. The configuration files are locted at install/pingfederate/server/default/conf

Configuration file and service implementation

Configuration file	RPC-based service implementation
cluster-account- locking.conf	Account Locking Service on page 200
cluster- artifact.conf	Artifact-Message Persistence and Retrieval Service on page 198
cluster- assertion-replay- prevention.conf	Assertion Replay Prevention Service on page 197
cluster- idp-session- registry.conf	IdP Session Registry Service on page 196
cluster-inter- request-state.conf	Inter-Request State-Management (IRSM) Service on page 195
cluster-session- revocation.conf	Back-Channel Session Revocation Service on page 199
cluster- sp-session- registry.conf	SP Session Registry Service on page 196

Administrator's Reference Guide

This guide provides information about using PingFederate to deploy a secure Internet single sign-on (SSO) solution based on the latest security and e-business standards.

Attribute mapping expressions

For new installations of PingFederate 10.1, the use of expressions is enabled by default. PingFederate 10.1 also provides a new Expression Admin administrative role, which allows you to map user attributes by using OGNL expressions.

If you have upgraded to PingFederate 10.1 from a previous version, the use of expressions is still enabled or disabled based on the configuration in the earlier version. Also, when upgrading PingFederate to 10.1 or later, administrative users who were granted the Admin role in the earlier installation are granted the Expression Admin role automatically.

Enabling and disabling expressions

As of PingFederate 10.1, the use of expressions is enabled by default. You can manually disable the use of expressions by editing a configuration file.

About this task

When upgrading PingFederate to 10.1 or later, administrative users who were granted the Admin role in the earlier installation are granted the Expression Admin role automatically.

You can disable the use of expressions by setting evaluateExpressions to false as described in the following procedure. Also, go to System # Server # Administrative Accounts and remove the Expression Admin role from all Admin users. Doing this will prevent Admin users from entering expressions into PingFederate if the evaluateExpressions element is set to true at a later time. For more information, see *Administrative accounts* on page 842.



Important:

If the current configuration contains expressions, disabling the feature causes errors during runtime processing.

Steps

1. Edit the org.sourceid.common.ExpressionManager.xml file, located in the <pf install>/ pingfederate/server/default/data/config-store directory.



Note:

If you have a clustered PingFederate environment, edit the configuration file on the console node.

2. Change the value of the element named evaluate Expressions to either true or false and save the file.

```
<?xml version="1.0" encoding="UTF-8"?>
<config xmlns="http://www.sourceid.org/2004/05/config">
   <item name="evaluateExpressions">true</item>
</config>
```



Note:

The absence of an installed default value does not necessarily disable the use of expressions. You can successfully import configuration archives containing expressions to facilitate backward compatibility when no value is present, and further use of the feature is enabled. The term "silent" is used for this condition in the server log.

3. If you have a stand-alone PingFederate environment, start or restart PingFederate.



If you are enabling expressions to use for mapping outbound provisioning attributes, you do not need to restart the PingFederate server.

- 4. If you have a clustered PingFederate environment:
 - a. Sign on to the PingFederate administrative console.
 - b. From System # Server # Cluster Management, click Replicate Configuration.

Result

When you enable expressions, they are available for use in multiple locations:

- The Source list under each of the administrative-console contract fulfillment windows
- The Show Advanced Criteria section on the Issuance Criteria window following each of the administrative-console contract fulfillment windows
- The provisioning attribute-mapping window when the Outbound Provisioning protocol is enabled

Construct OGNL expressions

Use OGNL expressions and syntax to evaluate and manipulate attribute values and return information based on the results.

OGNL is based on the Java programming language. You can transform a range of values into a text description or do the same for a sequence of ranges.

Use the # symbol to reference OGNL variables. For an identity provider (IdP), PingFederate provides predefined OGNL variables for IdP-adapter attributes, any attributes retrieved from datastores, and attributes for token authorization. For a service provider (SP), variables are available for attributes received in an assertion, an attribute query, and attributes for token authorization. For example, you can retrieve the SAML SUBJECT value with #SAML SUBJECT.



Note:

Use the following construction for any attributes from any source that contain special characters that cannot be parsed by OGNL: #this.get("<attribute name>").



Because OGNL uses the "at" symbol (@) to reference static Java methods, expressions containing the symbol must be enclosed in double quotes. Otherwise, expression parsing fails. For example, use #SAML SUBJECT="usr@msn.com", not #SAML SUBJECT=usr@msn.com.

Data store syntax

For datastore attributes with an attribute source ID, use the #this.get("ds.attr-sourceid.attribute name") syntax.

For datastore attributes without an attribute source ID, use the #this.get("ds.attribute name") syntax.

Other variable syntax

To access mapped attributes, use the #this.get("mapped.attribute name") syntax.

To access most context attributes, use the #this.get("context.attribute_name") syntax.

To access the HTTP Request context attribute, use the

#this.get("context.HttpRequest").getObjectValue() syntax.



Note:

The returned value is an instance of javax.servlet.http.HttpServletRequest. See http:// docs.oracle.com/javaee/7/api/javax/servlet/http/HttpServletReguest.html.

Sample OGNL expressions

OGNL expressions provide the flexibility to evaluate and manipulate values. These applications include using the following expressions to determine net worth, form a single sign-on (SSO) token, verify a user's group, retrieve a value from an HTTP request object, and check the authenticity of a client certificate..

General

In this sample expression, the value of the attribute "net-worth" is transformed first to eliminate any dollar signs or commas, then the result is evaluated to determine whether the user's net worth falls into a "bronze," "silver," or "gold" category.

```
#result=#this.get("net-worth").toString(),
#result=#result.replace("$",""),
#result=#result.replace(",",""),
#result < 500000 ? "bronze" :
#result < 1000000 ? "silver" : "gold"</pre>
```

Multivalued attribute

```
new org.sourceid.saml20.adapter.attribute.AttributeValue( {"Blue", "Gray",
    "Pink"})
```

This expression formulates a multivalued attribute in an SSO token.

```
<saml:Attribute Name="clrs" ...>
  <saml:AttributeValue ...>Blue</saml:AttributeValue>
  <saml:AttributeValue ...>Gray</saml:AttributeValue>
  <saml:AttributeValue ...>Pink</saml:AttributeValue>
</saml:Attribute>
```

and

```
{
    ...,
    "clrs": [
        "Blue",
        "Gray",
        "Pink"
    ],
    ...
}
```

In these truncated samples, clrs is the multivalued attribute. The former is a SAML assertion through a SAML service provider (SP) connection. The latter is a JSON web token (JWT) through a WS-Federation SP connection using JWT as the token type.

Token authorization

This expression verifies whether a user is a member of the "Engineering" or "Marketing" group.

```
#this.get("ds.memberOf").hasValue("CN=Mkt,OU=M,DC=contoso,DC=com")
    #this.get("context.VirtualServerId").toString().equals("Marketing")
):false
```

The following expression extracts the domain information out of an email address (mail) and returns true if it matches a specific domain.

```
#this.get("mail")!=null?
  #email=#this.get("mail").toString(),
  #atSign="@",
  #at=#mail.indexOf(#atSign),
  #at > 0?
      #domain=#mail.subject(#at+1),
      #domain.matches("(?i)example.com")
    ):false
):false
```

Note:

Line breaks are inserted to both samples for readability only. You must enter statements calling methods whose arguments are enclosed in quote on a single line.

This sample expression returns true when the IP address of the client is within the specified CIDR range of fe80::74da:14b:76d1:eba3/128.

```
#isWithinCidrRange =
@com.pingidentity.sdk.CIDROperations@isInRange(#this.get("context.ClientIp"), "fe80::740
```

The isInRange method supports both IPv4 and IPv6 CIDR notations.

HTTP request context

You can use the following example to retrieve a value from an HTTP request object. The expression retrieves the User-Agent HTTP header value and compares it against a value required for token authorization.

```
#this.get("context.HttpRequest").getObjectValue().getHeader("User-
Agent") .equals ("somevalue")
```

STS client authentication context

This security token service (STS) SSL Client Certificate Chain example checks that the issuer of the client certificate matches the specified distinguished name (DN).

```
#this.get("context.StsSSLClientCertChain").getObjectValue()
[1].getSubjectX500Principal().equals(new
javax.security.auth.x500.X500Principal("CN=Ping Identity
Engineering,OU=Engineering,O=Ping Identity,L=Denver,ST=CO,C=USA"))
```



Note:

#this.get("context.StsSSLClientCertChain").getObjectValue() returns an array of java.security.cert.X509Certificate instances. This array starts with the client certificate itself. For more information, see https://docs.oracle.com/javase/8/docs/api/java/security/cert/X509Certificate.html.

Issuance criteria and multiple virtual server IDs

Virtual server IDs offer critical information and functionality in the context of connections.

When you use virtual server IDs to connect to multiple environments in one connection, verifying at runtime the virtual server ID in conjunction with other end-user attributes, such as group membership, protects against unauthorized access.

For instance, both the sales and the support departments of contoso.com, the identity provider (IdP), have their own departmental subdomains, sales.contoso.com and support.contoso.com. The service provider (SP) identifies both environments under the parent domain, contoso.com.

In this scenario, you can configure the PingFederate IdP server to include both sales.contoso.com and support.contoso.com as the virtual server IDs in the SP connection.

If you use one IdP adapter to authenticate end users from both departments, use an OGNL expression to cross-check the virtual server ID information in the request and the end user's group membership information.

```
#this.get("ds.memberOf")!=null?
    #this.get("ds.memberOf").toString().matches("(?
i) CN=Eng, OU=E, DC=contoso, DC=com")
    #this.get("context.VirtualServerId").toString() == "Engineering"
  ) | |
    #this.get("ds.memberOf").toString().matches("(?
i) CN=Mkt, OU=M, DC=contoso, DC=com")
    #this.get("context.VirtualServerId").toString() == "Marketing"
):false
```

Note:

Line breaks are inserted for readability only. You must enter statements calling methods whose arguments are enclosed in quotes on a single line.

Expressions for OAuth and OpenID Connect uses cases

You can use OGNL expressions to retrieve various request-attributes through the HTTP Request Java object.

Client authentication method

The following sample expression retrieves the authentication method that a client uses. This sample expression is applicable to all clients.

```
#this.get("context.HttpRequest").getObjectValue().getAttribute("com.pingidentity.oauth.getObjectValue().getAttribute("com.pingidentity.oauth.getObjectValue().getObjectValue().getAttribute("com.pingidentity.oauth.getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue().getObjectValue(
```

Private key JSON web token (JWT)

In the following sample expressions, the former retrieves a claim value from the private key JWT with which a client authenticates and the latter retrieves the private key JWT itself. They are only applicable to clients using the private_key_jwt authentication method.

Retrieving the aud claim value

```
#claims =
#this.get("context.HttpRequest").getObjectValue().getAttribute("com.pingidentity.o
#claims.get("aud")
```

Retrieving the entire private key JWT

```
#this.get("context.HttpRequest").getObjectValue().getParameter("client assertion")
```

Using the OGNL edit window

Access the in-line OGNL editor and test expressions.

About this task

An in-line editor is available for OGNL expressions. The editor validates the expression and allows an administrator to enter input values and test the resulting output.

Steps

 To reach the OGNL editor, click Edit under Actions for an expression on any of the Attribute Fulfillment windows or click Test in the Show Advanced Criteria section on the Issuance Criteria window.



Note:

The test function does not work for the context. httpRequest attribute because its value is an object rather than text.

- To test an expression:
 - a. Enter an input value in the Value text box associated with the attribute.
 - b. Click the **Test**link near the bottom-right of the window.

If the expression contains no errors, the result appears under **Test Results**.



If you want to save changes to an expression, click **Update** under **Actions**. To discard changes, click the Cancel link under Actions. Click the Cancel button near the bottom of the window to discard all changes you made in the current task.

Authentication policies

Authentication policies are an optional configuration in PingFederate and help administrators implement complex authentication requirements.

As needed, administrators can configure one or more authentication selector instances to evaluate conditions of the requests and define policies to route the request to a series of approved authentication sources or deny the request based on the results from the authentication selector instances, authentication sources, or both. Administrators can also reuse an authentication policy by ending it with an authentication policy contract or a local identity profile and then applying the authentication policy contract in multiple use cases.

Authentication selectors provide a plugin capability for PingFederate to evaluate various conditions related to the requests. PingFederate comes bundled with a set of authentication selectors.

As an example, you can create an HTTP Header Authentication Selector to detect mobile browsers, a CIDR Authentication Selector to evaluate whether the users' IP addresses fall within your internal network ranges, or an HTTP Request Parameter Authentication Selector to identify identity provider (IdP) connections based on the PartnerIdpId parameter values provided in the service provider (SP)-initiated SSO requests.

Alternatively, you can create custom authentication selectors that suit your needs by using the PingFederate SDK.



Tip:

The Javadoc for PingFederate is located in the <pf install>/pingfederate/sdk/doc directory.

Managing authentication selector instances

You can manage authentication selectors on the **Selectors** window in the PingFederate administrative console.

Steps

- Go to Authentication # Policies # Selectors.
- To configure a new instance, on the Selectors window, click Create New Instance.
- To modify an existing instance, select it by its name under Instance Name.
- To remove an existing instance or to cancel the removal request, click **Delete** or **Undelete**.



Note:

You can only remove a selector instance if it is not deployed in any authentication policy.

Choosing a selector type

Choose an authentication selector instance from the list of available selector types.

Steps

- 1. Go to Authentication # Policies # Selectors.
- 2. Click Create New Instance.
- 3. In the **Instance Name** field, enter an instance name.
- 4. In the **Instance ID** field, enter an instance ID.
- 5. From the **Type**list, select the desired type of authentication selector.

Configuring an authentication selector instance

The configuration of an authentication selector instance varies depending on the authentication selectors deployed on your server.

- 1. Refer to subsequent topics for configuration steps of each of the bundled authentication selectors.
- 2. Complete the configuration.
 - a. On the **Summary** tab, click **Done**.
 - b. On the **Selectors** window, click **Save**.

Configuring the CIDR Authentication Selector

The CIDR Authentication Selector enables PingFederate to choose configured authentication sources or other selectors based on the IP address of an incoming single sign-on request.

About this task

Use this selector in authentication policies to choose from authentication sources that share a similar level of assurance, such as among multiple HTML Form Adapter instances or between a Kerberos Adapter instance and an X.509 identity provider (IdP) Adapter instance. For example, use this selector in authentication policies to route internal requests to a Kerberos Adapter instance.

Steps

- 1. Go to Authentication # Policies # Selectors to open the Selectors window.
- 2. On the Selectors window, click Create New Instance to start the Create Authentication Selector Instance workflow.
- 3. On the **Type** tab, configure the basics of this authentication selector instance.
- 4. On the Authentication Selector tab, click Add a new row to 'Networks' and enter a network range. Click **Update**.



To see the Add a new row to 'Networks' option, ensure you have set the Authentication Selector **Instance** type to **CIDR Authentication Selector** on the **Type** tab.

Example:

Sample IPv4 network range

Enter 192.168.101.0/24 to cover 256 IPv4 addresses, ranging from 192.168.101.0 through 192.168.101.255.

Sample IPv6 network range

Enter 2001:db8::/123 to cover 32 IPv6 addresses, ranging from 2001:db8:: through 2001:db8::1f.

5. Optional: Repeat the previous step to add more network ranges.



Note:

Display order does not matter.



If you want to include all IPv4 addresses for testing, add two separate ranges: 0.0.0.0/1 and 128.0.0.0/1. The CIDR Authentication Selector interprets a specification of 0.0.0.0/0 as an empty range rather than as a wildcard for all addresses.

Click Edit, Update, or Cancel to make or undo a change to an existing entry. Click Delete or **Undelete** to remove an existing entry or cancel the removal request.

6. Optional: Enter a Result Attribute Name value.



Note:

This field provides a means to indicate in the SAML assertion whether a network range was matched during processing; the value is either Yes or No. Any authentication sources configured as a result of this authentication selector must have their attribute contract extended with the value of the **Result Attribute Name** field in order to use its value to fulfill an attribute contract or for issuance criteria.

- 7. Complete the configuration.
 - a. On the **Summary** tab, click **Done**.
 - b. On the **Selectors** window, click **Save**.

Result

When you place this selector instance as a checkpoint in an authentication policy, it forms two policy paths: **Yes** and **No**. If the IP address of an incoming single sign-on (SSO) request matches one of the defined network ranges, the selector returns true. The policy engine regains control of the request and proceeds with the policy path configured for the result value of **Yes**. If the IP address of an incoming SSO request matches none of the defined network ranges, the selector returns false. The policy engine regains control of the request and proceeds with the policy path configured for the result value of **No**.

Configuring the Cluster Node Authentication Selector

The Cluster Node Authentication Selector enables PingFederate to choose configured authentication sources or other selectors based on the PingFederate cluster node that is servicing the request in authentication policies.

About this task

For example, this selector allows you to choose whether Integrated Windows Authentication (IWA) is attempted based on the PingFederate cluster node with which a Key Distribution Center (KDC) is associated.

Steps

- 1. Go to Authentication # Policies # Selectors to open the Selectors window.
- 2. On the **Selectors** window, click **Create New Instance** to start the **Create Authentication Selector Instance** workflow.
- 3. On the **Type** tab, configure the basics of this authentication selector instance.
- 4. On the **Authentication Selector** window, select the **Field Value** on which to branch policy paths. The authentication selector provides a means of choosing authentication sources at runtime based on the cluster node on which it is executing.

Node Index

Select Node Index to use the pf.cluster.node.index value specified in run.properties.

Node Tag

Select Node Tag to use the node.tags values specified in run.properties.

5. On the **Selector Result Values** window, specify the relevant node index or node tag values.



Each selector result value forms a policy path when you place this selector instance as a checkpoint in an authentication policy.

- a. In the Result Values field, enter a node index or node tag value based on your cluster configuration and click Add. This value should correspond to a node index or node tag of one of the engine nodes in the cluster.
- b. Optional: Add more values to differentiate criteria for authentication selection.



Display order does not matter.

Use the Edit, Update, and Cancel workflow to make or undo a change to an existing entry. Click **Delete** to remove an entry.

- 6. Complete the configuration.
 - a. On the Summary tab, click Done.
 - b. On the **Selectors** window, click **Save**.

Configuring the Connection Set Authentication Selector

The Connection Set Authentication Selector enables PingFederate to choose configured authentication sources or other selectors based on a match found between the target service provider (SP) connection used in a single sign-on (SSO) request and SP connections configured within PingFederate.

About this task

This selector allows you to override connection authentication selection on an individual connection basis in one or more authentication policies.

Steps

- 1. Go to Authentication # Policies # Selectors to open the Selectors window.
- 2. On the Selectors window, click Create New Instance to start the Create Authentication Selector **Instance** workflow.
- 3. On the **Type** tab, configure the basics of this authentication selector instance.
- 4. From the Type list, make sure you select Connection Set Authentication Selector.
- 5. Click Next. In the Authentication Selector window, click Add a new row to 'Connections'.
- 6. From the **Connection** list, select an SP connection and click **Update**.
- 7. Optional: Repeat the previous step to add more connections. Display order does not matter.

Click Edit, Update, or Cancel to make or undo a change to an existing entry. Click Delete or **Undelete** to remove an existing entry or cancel the removal request.

- 8. Complete the configuration.
 - a. On the **Summary** tab, click **Done**.
 - b. On the **Selectors** window, click **Save**.

Result

When you place this selector instance as a checkpoint in an authentication policy, it forms two Yes and No policy paths. If the invoking SP connection matches one of the connections from the set, the selector returns true. The policy engine regains control of the request and proceeds with the policy path configured for the result value of Yes. If the invoking SP connection matches none of the connections from the set, the selector returns false. The policy engine regains control of the request and proceeds with the policy path configured for the result value of No.

Configuring the Extended Property Authentication Selector

Configure this selector using values from the invoking browser-based single sign-on (SSO) or OAuth client.

About this task

The Extended Property Authentication Selector enables PingFederate to choose configured authentication sources or other selectors based on a match found between a selector result value and an extended property value from the invoking browser-based SSO connections or OAuth client.

Steps

- 1. Go to Authentication # Policies # Selectors to open the Selectors window.
- 2. On the Selectors window, click Create New Instance to start the Create Authentication Selector **Instance** workflow.
- 3. On the **Type** tab, configure the basics of this authentication selector instance.
- 4. On the Authentication Selector tab, select a property from the Extended Property list.



The extended property is the property that this selector instance should look for from the invoking connection or client, and compare the populated property value, or values if it is a multivalued extended property, against the selector result values defined in this selector instance.

- 5. On the **Selector Result Values** tab, specify one or more expected result values.
 - a. Enter the exact, case-sensitive, value under Result Values and click Add.
 - b. Optional: Add more values to differentiate criteria for authentication selection.

Display order might matter.

Expected result values are always sorted alphabetically in ascending order here.

When you place this selector instance as a checkpoint in an authentication policy, each selector result value forms a policy path. The display order of the resulting policy paths matches the display order here, which may impact the policy outcome. When the policy engine reaches this selector instance, the selector starts from top to bottom. As soon as it finds a match, it exits and returns true. The matching mechanism varies, depending on the type of the extended property selected in step 4.

Matching mechanism for single-value extended properties

The selector compares the property value populated in the invoking connection or client against the configured selector result value. When multiple selector result values exist, the selector starts from the top. If the current selector result value is a case-sensitive exact match, it returns true and exits. Otherwise, it moves on to the next selector result value and tries again.

For example, assume this selector instance, named ExtProps, is configured with expected result values of Alpha, Bravo, and Charlie. The invoking connection is populated with an extended property value of Bravo, and this selector instance is placed as a checkpoint in an authentication policy as follows.

```
ExtProps
+--Alpha
   <policy path>
```

Given this setup, the selector returns true and exits when it reaches the second selector result value. The policy engine regains control of the request and proceeds with the policy path configured for the selector result value of Bravo.

Matching mechanism for multivalued extended properties

The selector compares the property values populated in the invoking connection or client against the configured selector result value. If any one of the property values from the invoking connection or client is a case-sensitive exact match, the selector returns true and exits. When multiple selector result values exist, the selector starts from the top. If the current selector result value is a case-sensitive exact match to any one of the property values from the invoking connection or client, it returns true and exits. Otherwise, it moves on to the next selector result value and tries again.

For example, assume the previous selector instance remains. The invoking connection is populated with extended property values of Alpha and Charlie, and this selector instance remains as a checkpoint in an authentication policy.

In this scenario, the selector returns true and exits when it reaches the first selector result value. The policy engine regains control of the request and proceeds with the policy path configured for the selector result value of Alpha. Even though Charlie, the expected selector result value, is also a case-sensitive exact match to Charlie, one of the property values from the invoking connection, because the selector has already exited and returned control to the policy engine when it reaches Alpha, the policy engine will never execute the policy path configured for the selector result value of Charlie.

Use the **Edit**, **Update**, and **Cancel** workflow to make or undo a change to an existing entry. Click **Delete** to remove an entry.

- 6. Complete the configuration.
 - a. On the **Summary** tab, click **Done**.
 - b. On the **Selectors** window, click **Save**.

Example

Example

- 1. Go to System # Server # Extended Properties.
- 2. On the **Extended Properties**window, define a multivalued extended property, and name it **configStatus**.
- 3. Create an SP connection with the following characteristics:
 - On the **Extended Properties** window, add two values for the **configStatus** extended property: DEV and TEST.
 - On the **Attribute Source Mapping** window, map an authentication policy contract to the service provider (SP) connection. The policy contract name is APC.
- 4. Create an instance of the Extended Property Authentication Selector with the following characteristics:
 - On the **Type** tab, name the selector instance ExProps.
 - On the Authentication Selector tab, select configStatus from the list.
 - On the Selector Result Values tab, enter DEV and TEST.

5. Create and activate the following identity provider (IdP) authentication policy.

```
ExtProps
+--DEV
| OpenToken
 +--Fail: Done
  +--Success: APC
+--TEST
  HTML
   +--Fail: Done
   +--Success: APC
```

Configure each APC to fulfill values obtained from its preceding adapter instance.

When processing SSO requests intended for this SP connection, because the policy engine is able to match one of the populated property values, DEV, from the SP connection to the first selector result value, also DEV, it will always invoke the OpenToken IdP Adapter instance based on the DEV policy path. The TEST policy path is never executed for this SP connection.

On the other hand, if you remove DEV, an extended property value, from the SP connection, the policy engine will route SSO requests intended for this SP connection to the HTML Form Adapter instance based on the TEST policy path. The DEV policy path is never executed for this SP connection.

Configuring the HTTP Header Authentication Selector

The HTTP Header Authentication Selector enables PingFederate to choose configured authentication sources or other selectors based on a match found in a specified HTTP header.

About this task

Use this selector in one or more authentication policies to choose from authentication sources that share a similar level of assurance, such as among multiple HTML Form Adapters or between a Kerberos Adapter and an X.509 Adapter. For example, use this selector to choose an authentication source based on the user's browser identified by the User-Agent HTTP header.



Important:

Do not use this selector to determine whether an authentication source with a higher level of assurance should be bypassed because HTTP request headers could potentially be forged.

Steps

- Go to Authentication # Policies # Selectors to open the Selectors window.
- 2. On the Selectors window, click Create New Instance to start the Create Authentication Selector **Instance** workflow.
- 3. On the **Type** tab, configure the basics of this authentication selector instance.
- 4. On the Authentication Selector tab, click Add a new row to 'Results'.
- 5. Enter an expression for use when inspecting the HTTP header value of the target HTTP header under Match Expression, and click Update.



Note:

Wildcard entries are allowed, such as *value*.



Click Edit, Update, or Cancel to make or undo a change to an existing entry. Click Delete or **Undelete** to remove an existing entry or cancel the removal request.

- 7. In the Header Name field, enter the type of HTTP header you want the selector to inspect. This field is not case-sensitive.
- 8. Optional: To disable case-sensitive matching between the HTTP header values from the requests and the Match Expression values specified on this window, clear the Case-Sensitive Matching check box.

The Case-Sensitive Matching check box is selected by default.

- 9. Complete the configuration.
 - a. On the **Summary** tab, click **Done**.
 - b. On the **Selectors** window, click **Save**.

Result

When you place this selector instance as a checkpoint in an authentication policy, it forms two policy paths: Yes and No. If the value of the specified HTTP header matches one of the configured values, the selector returns true. The policy engine regains control of the request and proceeds with the policy path configured for the result value of Yes. If the value of the specified HTTP header matches none of the configured values, the selector returns false. The policy engine regains control of the request and proceeds with the policy path configured for the result value of **No**.

Example

Example

To detect the most common browsers based on the User-Agent HTTP request header, configure an HTTP Header Authentication Selector instance as follows.

1. Enter these entries under **Match Expression**.

Browser	Expression
Chrome	*Chrome*
Firefox	*Firefox*
Internet Explorer	*MSIE*
	Tip:
	For more information, see <i>User-agent string changes</i> from Microsoft.
Safari	*Safari*

2. In the **Header Name** field, enter User-Agent.

Configuring the HTTP Request Parameter Authentication Selector

The HTTP Request Parameter Authentication Selector enables PingFederate to choose configured authentication sources or other selectors based on query parameter values.

About this task

Use this selector in one or more authentication policies to choose from authentication sources that share a similar level of assurance, such as among multiple instances of the HTML Form Adapter or between a Kerberos Adapter instance and an X.509 Adapter instance. For example, use an instance of this selector to choose an authentication experience based on the reward program information indicated by a query parameter in the single sign-on (SSO) request.



Important:

Do not use this selector to determine whether an authentication source with a higher level of assurance should be bypassed because query parameters could potentially be forged.

Steps

- 1. Go to Authentication # Policies # Selectors to open the Selectors window.
- 2. On the Selectors window, click Create New Instance to start the Create Authentication Selector Instance workflow.
- 3. On the **Type** tab, configure the basics of this authentication selector instance.
- 4. On the **Authentication Selector** tab, configure the applicable selector instance settings.
 - a. Enter the exact, case-sensitive name of the request parameter in the HTTP Request Parameter Name field.



Important:

The policy engine is capable of tracking HTTP request parameters that it receives from the initial request and making them available to selector instances throughout the policy. If you plan on using this selector instance as the second, or subsequent, checkpoint in at least one

authentication policy, add the HTTP Request Parameter Name value on the Tracked HTTP Parameters window. For more information, see *Defining authentication policies* on page 242.

b. Optional: To disable case-sensitive matching between the HTTP request parameter values from the requests and the Match Expression values specified on the Selector Result Values window, clear the Case-Sensitive Matching check box.



The Case-Sensitive Matching check box is selected by default.

c. Optional: Enable policy paths to handle additional scenarios. For more information, see the following table.

Field	Description
Enable 'Any' Result Value	Each configured selector result value forms a separate authentication policy path.
	Select this check box if you want to enable a single policy path for the scenario where the HTTP request parameter value matches any one of the configured selector result values.
	This check box is not selected by default.
Enable 'No Match' Result Value	Selector evaluation fails and the next applicable authentication policy is executed when the HTTP request parameter value does not match any of the configured selector result values.
	Select this check box if you want to enable a policy path to handle this scenario.
	This check box is not selected by default.
Enable 'Not in Request' Result Value	Selector evaluation fails and the next applicable authentication policy is executed if the HTTP request parameter is not found.
	Select this check box if you want to enable a policy path to handle this scenario.
	This check box is not selected by default.

5. On the Selector Result Values window, enter a request parameter value under Result value, and then click Add.



Wildcard entries are allowed, such as *value*.



Important:

A more specific match is a better match and an exact match is the best match.

6. Optional: Repeat the previous step to add more request parameter values. Display order does not matter.



Note:

If you have not enabled the **Any** policy path in step 4c, each selector result value forms a policy path when you place this selector instance as a checkpoint in an authentication policy.

If you have enabled the **Any** policy path, only one policy path is formed.

Use the Edit, Update, and Cancel workflow to make or undo a change to an existing entry. Click **Delete** to remove an entry.

- 7. Complete the configuration.
 - a. On the **Summary** tab, click **Done**.
 - b. On the Selectors window, click Save.

Example

Example

Suppose you enter three selector result values, Central, Eastern, and Southern, on the Selector **Result Values** window, as illustrated in the following screen capture.

If you have not enabled any additional policy paths in step 4c, as you place this selector instance as a checkpoint in an authentication policy, three policy paths are extended from the selector instance, one for each of the configured selector result values.

Configuring the OAuth Client Set Authentication Selector

This selector allows you to override client authentication select on an individual client basis in one or more authentication policies.

About this task

The OAuth Client Set Authentication Selector enables PingFederate to choose configured authentication sources or other selectors based on a match found between the client information in an OAuth request and the OAuth clients configured in the PingFederate OAuth authorization server (AS).



Note:

The OAuth Client Set Authentication Selector is only applicable to OAuth clients using the authorization code or implicit flow.

- 1. Go to Authentication # Policies # Selectors to open the Selectors window.
- 2. On the Selectors window, click Create New Instance to start the Create Authentication Selector **Instance** workflow.
- 3. On the **Type** tab, configure the basics of this authentication selector instance.



If you do not see Add a new row to 'Clients', go back to the Type tab and ensure you have selected OAuth Client Set Authentication Selector from the Type list.

- 5. From the Client ID list, select an OAuth client and click Update.
- 6. Optional: Repeat the previous step to add more clients.

Display order does not matter.

Click Edit, Update, or Cancel to make or undo a change to an existing entry. Click Delete or **Undelete** to remove an existing entry or cancel the removal request.

- 7. Complete the configuration.
 - a. On the **Summary** tab, click **Done**.
 - b. On the **Selectors** window, click **Save**.

Result

When you place this selector instance as a checkpoint in an authentication policy, it forms two policy paths: Yes and No. If the invoking client matches one of the clients from the set, the selector returns true. The policy engine regains control of the request and proceeds with the policy path configured for the result value of Yes. If the invoking client matches none of the clients from the set, the selector returns false. The policy engine regains control of the request and proceeds with the policy path configured for the result value of No.

Configuring the OAuth Scope Authentication Selector

The OAuth Scope Authentication Selector enables PingFederate to choose configured authentication sources or other selectors based on a match found between the scopes of an OAuth authorization request and scopes configured in the PingFederate OAuth authorization server (AS).

Before you begin

Go to System # OAuth Settings # Authorization Server Settings and configure one or more scopes.

About this task

This selector allows you to control the strength of authentication based on client access requirements. For example, if a client requires write access to a resource, you can deploy an instance of the OAuth Scope Authentication Selector in one or more authentication policies to choose an adapter that offers a stronger form of authentication, such as the X.509 client certificate, instead of username and password.

- 1. Go to Authentication # Policies # Selectors to open the Selectors window.
- 2. On the Selectors window, click Create New Instance to start the Create Authentication Selector **Instance** workflow.
- 3. On the **Type** tab, configure the basics of this authentication selector instance.



Both common and exclusive scopes are available for selection.



Important:

This selector matches only scopes from OAuth authorization requests to the authorization endpoint, / as/authorization.oauth2. SAML single sign-on (SSO) requests do not match this authentication selector's criteria and result in a returned result value of No. If you are using this selector and selectors specific to SAML connections, list this selector first in the mapping list so that it takes precedence for OAuth without disrupting selector logic on SAML connections.

- 5. Complete the configuration.
 - a. On the Summary tab, click Done.
 - b. On the Selectors window, click Save.

Result

When you mark this selector instance as a checkpoint in an authentication policy, it forms two policy paths: Yes and No. If the requested scopes satisfy all the selected scopes, the selector returns true. The policy engine regains control of the request and proceeds with the policy path configured for the result value of Yes. If the requested scopes do not satisfy all the selected scopes, the selector returns false. The policy engine regains control of the request and proceeds with the policy path configured for the result value of No.

Configuring the Requested AuthN Context Authentication Selector

The Requested AuthN Context Authentication Selector enables PingFederate to choose configured authentication sources or other selectors.

About this task

This selector chooses authentication sources or selectors based on the authentication contexts requested by a service provider (SP) for browser single sign-on (SSO) requests, or a relying party (RP) for OAuth with OpenID Connect use cases in authentication policies.

For browser SSO, this authentication selector works in conjunction with SP connections with SAML 2.0 only, using the SP-initiated SSO profile. Other browser SSO protocols do not support authentication context. For OAuth, clients supporting the OpenID Connect protocol must include the optional acr values parameter in their authorization requests to indicate their preferred authentication context, or contexts.

- Go to Authentication # Policies # Selectors to open the Selectors window.
- 2. On the Selectors window, click Create New Instance to start the Create Authentication Selector Instance workflow.
- 3. On the **Type** tab, configure the basics of this authentication selector instance.

- 4. On the **Authentication Selector** tab, configure the applicable selector instance settings.
 - a. Select the Add or Update AuthN Context Attribute check box if you want to update the authentication context attribute value with the value specified in the Selector Result Values tab. Result:

When selected, which is the default, the check box on this window provides a means to:

- Add the value of the authentication context determined by the selector into the SAML assertion.
- When applicable, replace any value returned from the associated adapter instance with the selector-result value.
- b. Optional: Enable policy paths to handle additional scenarios.

For more information, refer to the following table.

Field	Description
Enable 'No Match' Result Value	Selector evaluation fails and the next applicable authentication policy is executed if the requested authentication context does not match any of the configured selector result values.
	Select this check box if you want to enable a policy path to handle this scenario. This check box is not selected by default.
Enable 'Not in Request' Result Value	Selector evaluation fails and the next applicable authentication policy is executed if no requested authentication context is found.
	Select this check box if you want to enable a policy path to handle this scenario. This check box is not selected by default.

- 5. On the **Selector Result Values** window, specify the authentication contexts to use as the criteria.
 - a. Enter the exact, case-sensitive parameter value under Result Values, and then click Add.



Note:

The value can include URIs defined in *Authentication Context for the OASIS Security Assertion Markup Language (SAML) 2.0* or any other value agreed upon with the partner.

b. Optional: Add more values to differentiate criteria for authentication selection.

Display order does not matter.

Each selector result value forms a policy path when you place this selector instance as a checkpoint in an authentication policy (regardless of whether you have enabled the **No Match** or **Not in Request** policy path in *step 4b*).

Use the **Edit**, **Update**, and **Cancel** workflow to make or undo a change to an existing entry. Click **Delete** to remove an entry.

- 6. Complete the configuration.
 - a. On the Summary tab, click Done.
 - b. On the Selectors window, click Save.

Configuring the Session Authentication Selector

You can configure the PingFederate Session Authentication Selector in the administrative console.

About this task

The Session Authentication Selector enables PingFederate to choose a policy path at runtime based on whether the user already has a PingFederate authentication session for a particular source.

- 1. Go to Authentication # Policies # Selectors to open the Selectors window.
- 2. On the Selectors window, click Create New Instance to start the Create Authentication Selector Instance workflow.
- 3. On the **Type** tab, configure the basics of this authentication selector instance.
- 4. On the Authentication Selector window, click Add a new row to 'Authentication Sources'.
- 5. Select an IdP adapter instance or an IdP connection from the list, enter a value under Result Value for the selected authentication source, then click **Update**.

The **Result Value** field controls the label shown for the policy path created by the selected authentication source.



Note:

You must enable authentication sessions for the selected authentication source, or globally for all authentication sources, on the Sessions window. Click Manage Sessions to review and configure authentication sessions.

6. Optional: Repeat the previous step to add more authentication sources.

Display order might matter.

When you place this selector instance as a checkpoint in an authentication policy, each selector result value forms a policy path. The display order of the resulting policy paths matches the display order here, which may impact the policy outcome. When the policy engine reaches this selector instance, the selector starts from top to bottom. It exits and returns true as soon as it finds a match.

As needed, use the up and down arrows to re-arrange the display order here, which also re-prioritizes the resulting policy paths.

In addition, when no session exists for any of the defined sources, the result value for the first authentication source is returned unless the Enable 'No Session' Result Value check box is selected, in which case an additional policy path is added as the last path when this selector instance is placed as a checkpoint in an authentication policy.

Click Edit, Update, or Cancel to make or undo a change to an existing entry. Click Delete or **Undelete** to remove an existing entry or cancel the removal request.

7. Optional: Select the Enable 'No Session' Result Value check box to create a separate policy path for the scenario where no session exists for any of the defined sources.

This check box is not selected by default.

- 8. Complete the configuration.
 - a. On the **Summary** tab, click **Done**.
 - b. On the **Selectors** window, click **Save**.

Result

When you place this selector instance as a checkpoint in an authentication policy, each selector result value forms a policy path that you can define the desired authentication experience and requirements.

Example

Example

The following screen capture illustrates a configuration where three authentication sources are defined and the Enable 'No Session' Result Value check box is selected.

Configuring a sample use case

Use the following sample setup to configure one of the common use cases where you have two categories of service providers (SPs).

Before you begin

For this sample use case, you must have the following components:

- An authentication policy contract
- Multiple SP connections. All connections use the same authentication policy contract as their sole authentication source
- Instances of the required adapters
- An instance of the Connection Set Authentication Selector to isolate high-value connections from the rest of the connections

About this task

The Session Authentication Selector enables PingFederate to choose a policy path at runtime based on whether the user already has a PingFederate authentication session for a particular source..

You need to enforce authentication requirements on two categories of service provider connections:

- For high-value connections, users must authenticate using the X.509 Adapter followed by the PingID Adapter.
- For low-value connections, users can authenticate using the HTML Form Adapter or the X.509 Adapter followed by the PingID Adapter.

To fulfill this use case, follow these configuration steps.

- 1. Go to Authentication # Policies # Selectors.
- 2. Create an instance of the Session Authentication Selector to account for authentication sessions acceptable for low-value connections.
 - a. Click Create New Instance.
 - b. On the **Type** tab, enter a name (for example, Sessions for low-value connections) and an ID; then select **Session Authentication Selector** from the list.
 - c. On the **Authentication Selector** tab, leave the **Enable 'No Session' Result Value** check box clear; then configure the following authentication source-to-result value entries.

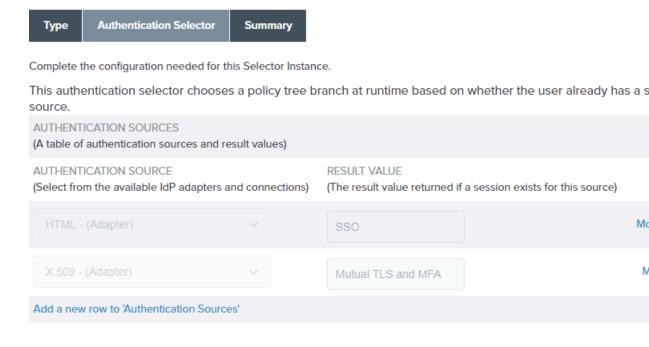
Authentication source (adapter instance name)	Result value (policy path label)
HTML	SSO

Authentication source (adapter instance name)	Result value (policy path label)
X.509	Mutual TLS and MFA

Example:

The following screen capture illustrates the

Manage Authentication Selector Instances | Create Authentication Selector Insta



Field Name	Field Value	Description
ENABLE 'NO SESSION' RESULT VALUE		By default, if no session above sources, the res source is returned. Selv return a separate resul

- d. On the Summary screen, click Done.
- e. On the **Manage Authentication Selector Instances** window, click **Save** to keep the newly configured authentication selector instance.
- 3. Go to Authentication # Policies # Policies.

- 4. On the **Policies** window, define an authentication policy for high-value connections.
 - a. Click Add Policy.
 - b. In the Name field, enter a name for the policy, such as High-value connections.
 - c. From the Policy list, select the instance of the Connect Set Authentication Selector that isolates high-value connections from the rest.
 - d. For the **No** policy path, select **Continue**.
 - e. For the **Yes** policy path, select the X.509 Adapter instance.
 - f. For the **X.509 Adapter instance** # Fail policy path, select **Done**.
 - g. For the X.509 Adapter instance # Success policy path, select the PingID Adapter instance.
 - h. Below the PingID Adapter instance, click Options.
 - i. On the Incoming User ID window, select the X.509 Adapter instance as the source and username as the attribute.



This step applies only to adapters that support a user identifier to be passed in from an earlier authentication source. The PingID Adapter requires this user identifier. For more information, see Specifying an incoming user ID on page 245.

- j. For the X.509 Adapter instance # Success # PingID Adapter instance # Fail policy path, select Done.
- k. For the X.509 Adapter instance # Success # PingID Adapter instance # Success policy path, select the authentication policy contract.
- I. Complete the contract mapping for the authentication policy contract. Example:

The following illustrates the policy created for high-value connections.

m. Click Done.

- 5. Define an authentication policy for low-value connections.
 - a. Click Add Policy.
 - b. Enter a name for the policy, such as Low-value connections.
 - c. From the Policy list, select the instance of the Session Authentication Selector. For more information, see step 2.
 - d. For the single sign-on (**SSO**) policy path, select the HTML Form Adapter instance.
 - e. For the HTML Form Adapter instance # Fail policy path, select Done.
 - f. For the HTML Form Adapter instance # Success policy path, select the authentication policy contract.
 - g. Complete the contract mapping for the authentication policy contract.
 - h. For the **Mutual TLS and MFA** policy path, select the X.509 Adapter instance.
 - i. For the **X.509** Adapter instance # Success policy path, select the PingID Adapter instance.
 - j. Below the PingID Adapter instance, click Options. Select the X.509 Adapter instance as the source and username as the attribute on the Incoming User ID window.



Z Tip:

This step only applies to adapters that support a user identifier to be passed in from an earlier authentication source. The PingID Adapter requires this user identifier. For more information, see Specifying an incoming user ID on page 245.

- k. For the X.509 Adapter instance # Success # PingID Adapter instance # Fail policy path, select Done.
- I. For the X.509 Adapter instance # Success # PingID Adapter instance # Success policy path, select the authentication policy contract.
- m. Complete the contract mapping for the authentication policy contract.

Example: The following illustrates the policy created for low-value connections.

- n. Click Done.
- o. To activate authentication polices for identity provider (IdP) browser SSO requests, adapter-toadapter requests, and browser-based OAuth authorization code and implicit flows, select the IdP Authentication Policies check box.

Example:

The following screen capture illustrates the policies created this sample use case.

6. To keep the newly configured authentication policies, click **Save**.

Policies

Administrators can create authentication policies to fulfill their authentication requirements.

An authentication policy is a tree of authentication sources, selector instances, or a combination of them, that defines the decision to route a request through a series of approved authentication sources with an optional authentication policy contract or a local identity profile at the end or to deny the request.

Administrators can enable authentication policies on identity provider (IdP) browser single sign-on (SSO) requests, adapter-to-adapter requests, and browser-based OAuth authorization code and implicit flows. Administrators can also enable authentication policies on service provider (SP)-initiated Browser SSO requests received at the /sp/startSSO.ping endpoint. Individual policies can be disabled, as needed.

The order of authentication policies matters because the policy engine starts from the first policy and works its way down. At runtime, the policy engine derives an authentication tree from the applicable policies and either approves or denies a request.

Policy paths, authentication policy contracts, and local identity profiles

Policy paths

An authentication policy starts with either a selector instance or an authentication source. Authentication sources and most selectors have two results, Success or Fail, Yes or No. Each result forms a policy path.

A policy path is open-ended if it contains only one or more selector instances, without any authentication sources. In this scenario, the policy engine continues to the next applicable authentication policy, if any.

A policy path is closed-ended if it contains one or more authentication sources, with or without any selector instances. A closed-ended path can optionally end with an authentication policy contract or a local identity profile.



Note:

A policy path is also closed-ended if it ends with an instance of a custom authentication selector that returns an IdP adapter instance ID or the connection ID of an IdP connection. Because the custom selector returns an authentication source, such a closed-ended path cannot end with an authentication policy contract or a local identity profile. Instead, it must end with an action of Done or Restart.

Authentication policy contracts and local identity profiles

An authentication policy contract can harness attribute values obtained from all authentication sources along the path leading up to it. Administrators can select the same authentication policy contract or local identity profile for different closed-ended paths, in one or more authentication policies, and fulfill them differently to suit the requirements. To enforce the same set of authentication policies in multiple use cases, map the authentication policy contract to the applicable Browser SSO connections and OAuth grant-mapping configuration.

A policy becomes more complex as the number of paths grows with the number of authentication sources and selector instances.

Multiple policies and runtime behavior

A complex policy can cover a lot of ground. However, depending on the authentication requirements, administrators can also create multiple policies to suit their needs.

When a request arrives at PingFederate, the policy engine skips all disabled policies and any closed-ended paths that are inapplicable to the request. A closed-ended path is considered inapplicable to a request in any of the following conditions:

- The local identity profile at the end of a path is associated with an authentication policy contract that is not mapped to the invoking use case or is blocked by the virtual server ID included in the request.
- The authentication policy contract at the end of a path is not mapped to the invoking use case or is blocked by the virtual server ID included in the request.
- The last authentication source at the end of a path, that does not end with an authentication policy contract or a local identity profile, is not mapped to the invoking use case or is blocked by the virtual server ID included in the request.



Virtual server IDs are not applicable to adapter-to-adapter mappings or OAuth use cases.

After pruning inapplicable policies and paths, the policy engine starts evaluating the request against the first applicable policy. Generally speaking, the policy engine moves on to the next applicable policy when it hits the end of an open-ended path, as indicated by an action of **Continue**, and stops when it hits the end of a closed-ended path, as indicated by an authentication policy contract or an action of **Done** or **Restart**. Depending on the policies, the policy engine might find an authentication source, a series of authentication sources, or no authentication source at all.

Default authentication sources

In the event that a request has only passed through an open-ended path and the policy engine finds no authentication source after evaluating the request through all the applicable policies, it picks the first applicable default authentication source. A default authentication source is considered applicable if it is mapped to the use case of the request.

If the policy engine cannot find a default authentication source and the Fail if policy engine finds no authentication source check box is not selected, PingFederate chooses an authentication source based on the following prioritized preferences:

1. If the request comes with an IdpAdapterId query parameter or a pfidpaid cookie, and if the authentication source specified by the query parameter or the cookie is mapped to the corresponding use case, PingFederate uses the specified authentication source. If the authentication source is not mapped, PingFederate denies the request and returns an error message.



Note:

If the request presents both the IdpAdapterId query parameter and the pfidpaid cookie, the IdpAdapterId query parameter takes precedence.

2. If the request comes with neither an IdpAdapterId query parameter nor a pfidpaid cookie, and if there is only one authentication source mapping, PingFederate uses the mapped authentication source.



Note:

If there are multiple authentication-source mappings, PingFederate returns the available authentication sources and lets the user authenticate through one of them. If the user selected the Remember selection check box and successfully authenticated, PingFederate returns a pfidpaid persistent cookie, identifying the user's preference.

If the Fail if policy engine finds no authentication source check box is selected, PingFederate denies the request and returns an error message.



Note:

If a request has passed through a closed-ended path, the policy engine has already found at least one authentication source for the user; in this scenario the policy engine ignores all default authentication sources.

Tracked HTTP request parameters

The policy engine is capable of tracking HTTP request parameters that it receives from the initial request and making them available to authentication sources, selector instances, and contract mappings throughout the policy.

Local identity profiles and authentication policy contracts

PingFederate empowers administrators to deliver a secure and easy-to-use customer authentication, registration, and profile management solution. A typical use case involves an HTML Form Adapter instance, a local identity profile, an authentication policy contract, and an IdP authentication policy. The HTML Form Adapter captures user attributes and maps them into an authentication policy contract through a local identity profile. In terms of configuration, the latter is accomplished by placing a local identity profile at the end of a policy path and completing the Local Identity Mapping # Contract Fulfillment configuration.

Defining authentication policies

Manage authentication policies and settings on the **Policies** window.

Steps

- 1. Go to Authentication # Policies # Policies.
- 2. On the Policies tab, select the IdP Authentication Policies check box if you want to enable authentication policies for identity provider (IdP) browser single sign-on (SSO) requests, adapter-toadapter requests, and browser-based OAuth authorization code and implicit flows.
- 3. Select the SP Authentication Policies check box if you want to enable authentication policies for service provider (SP)-initiated browser SSO requests received at the /sp/startSSO.ping endpoint.



Note:

Selecting the SP Authentication Policies check box does not enable authentication policies for IdP browser SSO requests, adapter-to-adapter requests, and browser-based OAuth authorization code and implicit flows.

- 4. Select the Fail if policy engine finds no authentication source check box if you want PingFederate to deny the requests and to return an error message when the policy engine finds no authentication source or authentication policy contract from the applicable policies and none of the default authentication sources are applicable. This check box is not selected by default.
- 5. On the **Policies** window, click **Add Policy** to create an authentication policy.



If you want to create a new policy based on an existing policy, select the **Copy** action.

- a. Enter a name and, optionally, a description of the policy.
- b. From the **Policy** list, choose an authentication source, an IdP adapter instance or an IdP connection, or a selector instance.



Note:

If you start this new policy by copying an existing policy, your new policy is pre-populated. Modify the policy to suit your new use cases.



When implementing your authentication requirements, think of authentication sources and selectors as checkpoints.

Options

For the PingID Adapter, IdP adapters developed using the

IdpAuthenticationAdapterV2 interface from the PingFederate SDK, including the HTML Form Adapter, and SAML 2.0 IdP connections supporting the SP-initiated browser SSO profile, you can specify a user ID to be passed in from an earlier-factor adapter. Click Options and follow the on-window instructions to select the source and the attribute to

be used as the incoming user ID.

Rules

For any authentication source, you can optionally create one or more rules to define additional successful results. For example, if you want to deploy multifactor authentication using the PingID Adapter in stages by groups, you can create a rule to check for group membership information and only apply the PingID authentication flow to users who are members of certain groups.

Click Rules and follow the on-window instructions to manage your rules.

All results, including those based on rules, are displayed under the selected authentication source or selector instance. Each result forms a policy path.

- c. For each policy path, select a policy action from the list.
 - If additional processing is required, repeat step 4b.
 - If the policy path is extended from an authentication source and it is the end of the path, select **Done** or **Restart**, which marks this policy path a closed-ended path.



Z Tip:

A policy path is closed-ended if it contains one or more authentication sources, with or without any selector instances. A closed-ended path can optionally end with an authentication policy contract or a local identity profile.

If you need to reuse an authentication policy in multiple use cases, select an authentication policy contract or a local identity profile as the last policy action of a path, configure its contract fulfillment, and map the authentication policy contract to the applicable Browser SSO connections or OAuth grant-mapping configuration. Click ... Mapping underneath your selection and then follow the on-window instructions to complete the contract fulfillment configuration.



Note:

A policy path is also closed-ended if it ends with an instance of a custom authentication selector that returns an IdP adapter instance ID or the connection ID of an IdP connection. Because the custom selector returns an authentication source, such a closed-ended path cannot end with an authentication policy contract or a local identity profile. Instead, it must end with an action of Done or Restart.

The **Restart** policy action provides users the opportunity to do over. When triggered, the policy engine routes the requests back to the first checkpoint of the invoked authentication policy. It makes most sense to use the **Restart** policy action for a **Fail** policy path if the policy engine can route the request differently based on user input prompted by an authentication

source. For a sample use case, see step 5k in Enabling third-party identity providers on page 341.



Undesirable looping behaviors can occur if you select **Restart** for the **Fail** path at the root of an authentication policy tree. PingFederate mitigates this risk by automatically limiting the number of policy restarts per transaction.

If the policy path is extended from a selector instance and it is the end of the path without any prior authentication source, select Continue, which leaves this path as an open-ended path.



Tip:

A policy path is open-ended if it contains only selector instances without any authentication sources. In this scenario, the policy engine continues to the next applicable authentication policy, if any.

d. Click Done to go back to the Policies window.

Result:

Your policy is enabled by default. As needed, toggle its status to disable the policy.

6. Optional: Repeat *step 4* to create additional authentication policies.



Important:

The order of authentication policies matters because the policy engine starts from the first policy and works its way down. As needed, reorder your policies by using the up and down arrows.

- 7. If any individual policy is no longer required, select the **Delete** action or toggle its status to disable the policy.
- 8. Optional: On the Policies # Default Authentication Sources tab, select one or more default authentication sources from the list for the policy engine to fall back on when it finds no authentication source from the applicable policies.



Important:

Order matters because the policy engine starts from the first default authentication source on the list and works its way down. As needed, reorder your authentication sources by using the up and down arrows. There is no default selection.

9. Optional: On the Policies # Tracked HTTP Parameters tab, add one or more HTTP request parameters to be tracked throughout a request.



Important:

For each instance of the HTTP Request Parameter Authentication Selector that you place in a policy as the second, or subsequent, checkpoint, add its configured HTTP Request Parameter Name value here. By doing so, the policy engine preserves the parameter it receives from the initial request and makes it available to the selector instance throughout the policy. For more information, see Configuring the HTTP Request Parameter Authentication Selector on page 229.

10. Click Save.

Specifying an incoming user ID

You can modify authentication policies to make use of a user identifier at request time.

About this task

Some authentication sources make use of a user identifier at request time. For example:

- The PingID Adapter requires a user ID to be passed in from an earlier-authentication step to perform multifactor authentication.
- The HTML Form Adapter and custom IdP adapters developed using the IdpAuthenticationAdapterV2 interface from the PingFederate SDK can pre-populate username information based on an incoming user ID.
- A SAML 2.0 identity provider (IdP) connection can use an incoming ID to specify the Subject value in its authentication requests.
- An OpenID Connect IdP connection can leverage an incoming user ID to specify a login hint parameter value in its OAuth authorization requests.

To address these use cases, use the **Options** # **Incoming User ID** dialog to specify the source and the attribute of the incoming user ID in an authentication policy.

You can select any IdP adapter instance or IdP connection that has been placed in the same policy path ahead of the current authentication source to be the source of the incoming user ID. After you select a source, choose an attribute from the selected IdP adapter contract or IdP connection. At runtime, the attribute value becomes the incoming user ID.

Alternatively, you can use the originating SAML 2.0, WS-Federation, or OpenID Connect authentication request as the source. In this scenario, the incoming user ID is derived from the Subject element in a SAML 2.0 authentication request, the username parameter in a WS-Federation authentication request, or the login hint parameter in an OpenID Connect authentication request.

Steps

- 1. Go to Authentication # Policies # Policies. On the Policies window, select the applicable authentication policy.
- 2. On the **Policy** tab, locate the authentication source that you need to provide an incoming user ID and then click **Options** underneath it.
- 3. On the Incoming User ID dialog, select the source of the incoming user ID from the Source list.



Note:

If you want the policy engine to derive the incoming user ID from the originating SAML 2.0 or WS-Federation authentication request, select **Context**.

4. Select an attribute of the incoming user ID from the **Attribute** list.



Note:

If you have selected Context in the previous step, select Requested User to derive the incoming user ID from the Subject element, the username parameter, or the login hint parameter in the SAML 2.0, WS-Federation, or OpenID Connect authentication request, respectively.

Result:

If a request does not originate from a SAML 2.0, WS-Federation, or OpenID Connect authentication request, or if the SAML 2.0, WS-Federation, or OpenID Connect authentication request does not include the optional Subject element, username parameter, or login hint parameter, the policy engine advances without providing username information to the authentication source.

5. On the **Incoming User ID** dialog window, click **Done**.

6. On the **Policy** tab, continue with the rest of your policy configuration.

Configuring rules in authentication policies

PingFederate supports more granular control through the use of rules in authentication policies.

About this task

An authentication source in an authentication policy has two results, Fail or Success, for which you can set one of the following actions:

- Append another authentication source for further processing.
- Append a selector for further processing.
- Select **Done** to terminate the authentication policy, making it a closed-ended path.
- Select an authentication policy contract or a local identity profile, also terminating the authentication policy, making it a closed-ended path.



Z Tip:

A policy path is closed-ended if it contains one or more authentication sources, with or without any selector instances. A closed-ended path can optionally end with an authentication policy contract or a local identity profile.



Note:

A policy path is also closed-ended if it ends with an instance of a custom authentication selector that returns an IdP adapter instance ID or the connection ID of an IdP connection. Because the custom selector returns an authentication source, such a closed-ended path cannot end with an authentication policy contract or a local identity profile. Instead, it must end with an action of **Done** or **Restart**.

By applying multiple rules to an authentication source, an administrator can define additional, successful, results based on attribute values from the authentication source and set different action for each result.

For example, your OpenToken IdP Adapter instance returns an attribute, EmployeeType, that identifies the employee profile; a value of temp indicates the user is a contractor. Your organization mandates that all contractors must authenticate successfully against the OpenToken identity provider (IdP) adapter, followed by another IdP adapter, such as an instance of the PingID Adapter for multifactor authentication. To fulfill this authentication requirement, you can define a successful result by adding a rule to evaluate the EmployeeType value, and then select the PingID Adapter instance as the action for this match.

When multiple rules exist for a given authentication source, the first match wins. If no rule returns a match, administrators have the option to treat the authentication as successful or failure.

- 1. Go to Authentication # Policies # Policies. On the Policies window, select the applicable authentication policy.
- 2. On the **Policy** tab, locate the authentication source that you want to define additional successful results for further processing, and then click **Rules** underneath it.
- 3. On the Rules window, select an attribute from the Attribute Name list.

The choices are:

- equal to
- equal to (case insensitive)
- equal to DN
- not equal to
- not equal to (case insensitive)
- not equal to DN
- multi-value contains
- multi-value contains (case insensitive)
- multi-value contains DN
- multi-value does not contain
- multi-value does not contain (case insensitive)
- multi-value does not contain DN

Use one of the first six choices only for attributes consisting of a single value. Use the multi-value conditions when you want PingFederate to verify whether an attribute contains or does not contain the specified value in its attribute list.



CAUTION:

Using a single-value condition when an attribute has multiple values causes the condition to evaluate consistently to false.

- 5. In the Value field, enter the desired value to be compared against the attribute value from the authentication source.
- 6. In the **Result** field, enter a unique label.
- 7. Optional: To add another rule, click **Add** and repeat steps 3 to 6.
- 8. If any individual rule is no longer required, select the **Delete** action.
- 9. Select the **Default to Success** check box if you want the policy engine to treat the authentication attempt as successful when no rules return a match.



Note:

The **Default to Success** check box is selected by default. When you clear this check box, the policy engine treats the attempt as a failure when no rules return a match.

10. To close the **Rules** window, click **Done**.

Result:

Your policy is now updated with a new policy path, or paths if you have added multiple rules.

For instance, if you have added two rules with labels Contractors, the first rule, and Senior executives, the second rule, to an authentication source, you should see the following results in the policy:

- Fail
- Contractors, a new result based on the first rule
- Senior executives, a new result based on the second rule
- Success, available only when the Default to Success check box is selected
- 11. On the **Policy** window, continue with the rest of your policy configuration.

PingFederate lets you configure authentication policies based on group membership information through the use of rules.

About this task

Assume you have created the following authentication policy to enforce multifactor authentication using PingID after the users have successfully authenticated against an HTML Form Adapter instance.

While this policy satisfies the authentication requirements, you might prefer to roll out multi-factor authentication based on group membership over a period of time. To accomplish this policy deployment strategy, you can use rules to define the applicable groups and set different policy actions accordingly.

As an example, suppose you want to enforce PingID multi-factor authentication to two Active Directory (AD) groups:

- CN=helpdesk,OU=IT,DC=example,DC=com (IT helpdesk personnel)
- CN=leads,OU=IT,DC=example,DC=com (Leaders in the IT department)

Steps

1. If you have not done so, create a new AD group, for example, CN=PingIDRequired, OU=IT, DC=example, DC=com, and place those two groups as members of the new group.



Tip:

Generally speaking, this step streamlines the process of deploying your authentication policies to additional groups of users later. In other words, when you are ready to roll out your authentication policies to more users, simply add the applicable groups as members of the new group. This way, you are not required to make any changes to the authentication policies, once they are configured.

- 2. If you have not done so, follow these steps to extend the HTML Form Adapter instance to return group membership information from your AD.
 - a. Extend the HTML Form Adapter instance with the memberof attribute on the Extended Contract
 - b. Configure PingFederate to fulfill the memberOf attribute from your AD.



Note:

Because the actual groups are nested inside the new group created in step 1, configure the IdP adapter contract to pull the memberOf attribute values with nested groups on the LDAP Directory Search window. For more information, see Defining the IdP adapter contract on page 386.



- 3. Go to Authentication # Policies # Policies. On the Policies window, select the applicable authentication policy.
- 4. On the **Policy** window, click **Rules** underneath the HTML Form Adapter instance.

- 5. In the Rules dialog, add a rule to check for membership information obtained from the HTML Form Adapter instance.
 - a. From the Attribute Name list, select memberOf.
 - b. Select how PingFederate should compare the value that you are going to specify in the next step against the attribute value from the HTML Form Adapter instance. For example, multi-value contains DN works well for the memberOf AD user attribute.
 - c. Enter the distinguished name (DN) of the new group created in step 1 in the Value field, such as CN=PingIDRequired, OU=IT, DC=example, DC=com.
 - d. In the Result field, enter a label, for example, PingID users.
 - e. Leave the **Default to Success** check box as selected.



When the **Default to Success** check box is selected, you can set the action for the scenario where none of the rules returns a match.

Example:

Your **Rules** dialog should be similar to the following sample.

- 6. To close the **Rules** window, click **Done**.
- 7. For the new HTML Form # PingID users policy path, select your PingID Adapter instance as the policy action from the list.



To open the **Incoming User ID** dialog, click **Options**, underneath the PingID Adapter instance. Configure the source of the user ID required by the PingID Adapter, and then click **Done** to close the dialog. For more information, see Specifying an incoming user ID on page 245.

Result:

Your policy should be similar to the following sample.

There are four policy paths:

- HTML Form # Fail
- HTML Form # PingID users # Fail
- HTML Form # PingID users # Success
- HTML Form # Success
- 8. Update your policy as follows:

HTML Form # Fail

Leave **Done** as the policy action. At runtime, PingFederate terminates the request and returns an error message to the user.

HTML Form # PingID users # Fail

Select **Done** as the policy action. At runtime, PingFederate terminates the request and returns an error message to the user.

HTML Form # PingID users # Success

HTML Form # Success

Reconfigure the policy action for this policy path. Select an authentication policy contract as the policy action. Click Contract Mapping to complete its fulfillment. At runtime, PingFederate fulfills the authentication policy contract and carries on with the request.

Result:

Your policy should be similar to the following sample.

- 9. To close the **Policy** window, click **Done**.
- 10. On the **Policies** window, click **Save**.



Note:

Group membership is only one of the possible factors that you can use to define additional policy paths and their policy actions. In general, you can use any attributes available from the authentication source when configuring rules.

Applying policy contracts or identity profiles to authentication policies

To apply an authentication policy contract to a policy, select an authentication policy contract or a local identity profile as the last action of one or more closed-ended paths and configure fulfillment for each contract.

About this task

An authentication policy contract can harness attribute values obtained from all authentication sources along the path leading up to it. Administrators can select the same authentication policy contract or local identity profile for different closed-ended paths, in one or more authentication policies, and fulfill them differently to suit the requirements. To enforce the same set of authentication policies in multiple use cases, map the authentication policy contract to the applicable Browser SSO connections and OAuth grant-mapping configuration.

Steps

- 1. Go to Authentication # Policies # Policies. On the Policies window, select the applicable authentication policy.
- 2. On the **Policy** window, locate all closed-ended paths in the policy.

A policy path is closed-ended if it contains one or more authentication sources, with or without any selector instances. A closed-ended path can optionally end with an authentication policy contract or a local identity profile.



Note:

A policy path is also closed-ended if it ends with an instance of a custom authentication selector that returns an IdP adapter instance ID or the connection ID of an IdP connection. Because the custom

Consider the following sample policy.

This policy has two selector instances, **Test** and **Retail**, two identity provider (IdP) adapter instances, and five policy paths:

- Test # No # HTML Form # Fail
- Test # No # HTML Form # Success # Retail # No
- Test # No # HTML Form # Success # Retail # Yes # PingID # Fail
- Test # No # HTML Form # Success # Retail # Yes # PingID # Success
- Test # Yes

The first four paths are closed-ended while the last path is open-ended.

- 3. Select **Done** as the policy action for the following paths:
 - Test # No # HTML Form # Fail
 - Test # No # HTML Form # Success # Retail # Yes # PingID # Fail

Result:

At runtime, PingFederate terminates the request and returns an error message to the user.

- 4. Select the applicable authentication policy contract or local identity profile as the policy action for the rest of the closed-ended paths:
 - Test # No # HTML Form # Success # Retail # No
 - Test # No # HTML Form # Success # Retail # Yes # PingID # Success

Suppose your use case does not involve consumer authentication, registration, and profile management. It makes sense to select an authentication policy contract for the **PingID** # **Success** result, because the users have successfully met all your authentication requirements.

At runtime, PingFederate fulfills the authentication policy contract and carries on with the request.

Depending on your use case, you might also select an authentication policy contract for the **PingID** # **Fail** result, possibly with an attribute indicating that the users have failed a certain part of your authentication requirements, and make other authorization decision using the Token Authorization framework in the applicable connections later.

- 5. For each selected authentication policy contract, if any, click **Contract Mapping** and then complete the **Manage Authentication Policies** # **Authentication Policy Contract Mapping** workflow to complete the configuration. For more information, see *Configuring contract mapping* on page 252.
- 6. For each selected local identity profile, if any, click **Local Identity Mapping** and then complete the **Manage Authentication Policies** # **Inbound Mapping & Contract Fulfillment** workflow to complete the configuration. For more information, see *Configuring local identity mapping* on page 252.
- 7. Select **Continue** as the policy action for the open-ended path **Test** # **Yes**.

Result

At runtime, PingFederate skips to the next policy. Your policy should be similar to the following sample.

- 8. To close the **Policy** window, click **Done**.
- 9. On the **Policies** window, click **Save**.

Configuring contract mapping

Configure authentication policy contract grant mappings in the PingFederate administrative console.

Steps

- 1. Optional: Go to Authentication # OAuth # Authentication Policy Contract Mapping.
- 2. Optional: From the Authentication Policy Contract list, click the desired mapping or select the desired mapping from the Authentication Policy Contract list.



Note:

If you do not have an authentication policy contract mapping already configured, go to Authentication # Policies # Policy Contracts and configure and save a new contract.

- 3. Optional: On the Attribute Sources & User Lookup window, click Add Attribute Source to configure datastore queries.
- 4. On the **Contract Fulfillment** tab, fulfill the selected contract.



Note:

If the selected closed-ended path contains more than one authentication source, you have access to attributes obtained successfully from the previous authentication sources along the same path.

For example, referring to the earlier policy in Applying policy contracts or identity profiles to authentication policies on page 250, if you select an authentication policy contract for the PingID (Adapter) # Success result, you can map attributes from the HTML Form Adapter and the PingID Adapter.

Besides the preceding identity provider (IdP) connection or IdP adapter instance, you can also use dynamic text, attribute mapping expression, if enabled, and tracked HTTP request parameter, if configured, as the source of fulfillment.

- 5. Optional: On the Issuance Criteria tab, configure conditions to be validated before issuing an authentication policy contract. For more information, see Defining issuance criteria for contract or local identity mapping on page 253.
- 6. On the **Summary** tab, review your configuration, modify as needed, and then click **Done**.
- 7. On the **Policy** window, continue with the rest of your policy configuration.

Configuring local identity mapping

You can configure your local identity mapping in the PingFederate administrative console.

Steps

1. On the **Inbound Mapping** tab, configure the attribute mappings for registration and profile management.



Note:

At runtime, PingFederate fulfills the value of the pf.local.identity.unique.id builtin local identity field based on this configuration and passes the value to PingDirectory. PingDirectory uses this value to determine whether such identity has already been created. The pf.local.identity.unique.id field value should therefore be mapped from the subject identifier of the preceding authentication source. You can also map other local identity fields so that PingFederate can streamline the registration process by pre-populating values on the registration page.



Note:

This configuration overrides the default field values configured within the local identity profile. For more information, see Configure a local identity field.

This tab does not apply and stays hidden if your use case does not involve registration and profile management. See Enabling third-party identity providers without registration on page 352.

- 2. Optional: On the Attribute Sources & User Lookup tab, click Add Attribute Source to configure datastore queries.
- 3. On the Contract Fulfillment tab, fulfill the authentication policy contract associated with the selected local identity profile.

If the selected closed-ended path contains more than one authentication source, you have access to attributes obtained successfully from the previous authentication sources along the same path.

For example, select your local identity profile under **Source** and the desired local identity field under Value.



Note:

If your use case does not involve registration or profile management, the source of fulfillment is limited to the preceding identity provider (IdP) connection or IdP adapter instance, dynamic text, attribute mapping expression, if enabled, and tracked HTTP request parameter, if configured.

- 4. Optional: On the Issuance Criteria tab, configure conditions to be validated before issuing an authentication policy contract. For more information, see Defining issuance criteria for contract or local identity mapping on page 253.
- 5. On the **Summary** tab, review your configuration, modify as needed, and then click **Done**.
- 6. On the **Policy** window, continue with the rest of your policy configuration.

Defining issuance criteria for contract or local identity mapping

You must define certain criteria for contract or local identity mapping in PingFederate to process a request.

About this task

On the Issuance Criteria tab, define the criteria to satisfy in order for PingFederate to further process a request. Use this token authorization feature to conditionally approve or reject requests based on individual attributes.

Begin this optional configuration by choosing the source that contains the attribute to verify. Some sources are common to almost all use cases, such as Mapped Attributes. Other sources depend on the type of configuration, such as JDBC. Irrelevant sources are automatically hidden. Once you select a source, choose the attribute to verify. Depending on the selected source, the available attributes or properties vary. Specify the comparison condition and the desired value to compare to.

You can define multiple criteria, which must all be satisfied in order for PingFederate to move a request to the next phase. A criterion is satisfied when the runtime value of the selected attribute matches or does not match the specified value, depending on the chosen comparison method. The multi-value contains ... or multi-value does not contain ... comparison methods are intended for attributes that can contain multiple values. Such a criterion is considered satisfied if one of the multiple values match or does not match the specified value. Values are compared verbatim. If you require complex evaluations, including conditional criteria or partial matching, define them using attribute mapping expressions.

Note:

All criteria defined must be satisfied, or evaluated as true, for a request to move forward, regardless of how the criteria were defined. As soon as one criterion fails, PingFederate rejects the request and returns an error message.

Steps

1. From the **Source** list, select the attribute's source.

Depending on the selection, the Attribute Name list populates with associated attributes. See the following table for more information.

Source	Description
Adapter	Select to evaluate attributes from any preceding IdP adapter instance.
IdP Connection	Select to evaluate attributes from any preceding IdP connection.
Local Identity	Select to evaluate any local identity fields.
	Not applicable for the Contract Mapping configuration.
JDBC, LDAP, or other types of datastore (if configured)	Select to evaluate attributes returned from a data source.
Mapped Attributes	Select to evaluate the mapped attributes.
Tracked HTTP	Select to evaluate tracked request parameters.
Parameters	Visible and applicable only if at least one HTTP request parameter has been configured on the Tracked HTTP Parameters tab of the Policies window. For more information, see .

- 2. From the **Attribute Name** list, select the attribute to be evaluated.
- 3. From the **Condition** list, select the comparison method.

Available methods:

- equal to
- equal to (case insensitive)
- equal to DN
- not equal to
- not equal to (case insensitive)
- not equal to DN
- multi-value contains
- multi-value contains (case insensitive)
- multi-value contains DN
- multi-value does not contain
- multi-value does not contain (case insensitive)
- multi-value does not contain DN



Note:

The first six conditions are intended for single-value attributes. Use one of the multi-value ... conditions for PingFederate to validate whether one of the attribute values matches the specified value. When an attribute has multiple values, using a single-value condition causes the criteria to fail.



Note:

Values are compared verbatim. If you require complex evaluations, including conditional criteria or partial matching, define them using attribute mapping expressions. For more information, see Attribute mapping expressions on page 214.

5. In the **Error Result** field, enter a custom error message.

To use localized descriptions, enter a unique alias in the Error Result field, such as someIssuanceCriterionFailed. Insert the same alias with the desired localized text in the applicable language resource files, located in the <pf install>/pingfederate/server/ default/conf/language-packs directory.

If not defined, PingFederate returns ACCESS DENIED when the criterion fails at runtime.

- 6. Click Add.
- 7. Optional: Repeat to add more criteria.
- 8. If you require complex evaluations, including conditional criteria or partial matching, define them using attribute mapping expressions. For more information, see Attribute mapping expressions on page 214.
 - a. Click Show Advanced Criteria.
 - b. In the **Expression** field, enter the required expressions.
 - c. Optional: In the Error Result field, enter an error code or message.



Note:

If the expressions resolve to a string value instead of true or false, the returned value overrides the Error Result field value.

- d. Click Add.
- Optional: Click Test, enter values in the applicable fields, and verify the results.
- f. Optional: Repeat to add multiple criteria using attribute mapping expressions.

Mapping a policy contract to multiple use cases

The last step to reuse an authentication policy in multiple service provider (SP) connections is to map the authentication policy contract into the applicable SP connections.

About this task

In general, for identity provider (IdP) browser single sign-on (SSO) use cases, if you select authentication policy contracts in your authentication policies then you must map the authentication policy contracts to the applicable SP connections.

Steps

- 1. Go to Applications # Integration # SP Connections.
- 2. Select the applicable SP connection from the list of connections.
- 3. On the Activation & Summary tab, click Authentication Source Mapping.
- 4. Click Map New Authentication Policy and use the in-product help on each screen as needed to map the authentication policy contract into the SP connection.

Result

Similarly, to reuse an authentication policy for browser-based OAuth authorization code and implicit flows, map the authentication policy contract to the applicable browser SSO connections and OAuth grantmapping configuration. For more information, see Managing authentication policy contract grant mapping on page 538.

SP authentication policies

Service provider (SP) authentication policies provide a means for you to impose authentication requirements on SP-initiated browser single sign-on (SSO) requests received at the /sp/ startSSO.ping endpoint.

When you enable this optional feature, you create policies that the PingFederate SP server can use to find the applicable SP adapter instance to access target applications. For this reason, you must configure the target applications to provide the SpSessionAuthnAdapterId parameter or the TargetResource parameter, or both, in their SP-initiated SSO requests.

If you prefer to provide the TargetResource parameter without the SpSessionAuthnAdapterId parameter, you must go to Applications # Integration # Target URL Mapping and configure entries to map the TargetResource values to the applicable SP adapter instances.



Note:

SP authentication policies only apply to SP-initiated browser SSO requests received at the /sp/ startSSO.ping on page 1139 SP application endpoint. They do not apply to unsolicited SSO requests received at the SP protocol endpoints.

In addition, enabling SP authentication policies does not enable authentication policies for identity provider (IdP) browser SSO requests, adapter-to-adapter requests, and browser-based OAuth authorization code and implicit flows.

For more information and configuration steps, see the subsequent sample use cases.

Configuring an SP authentication policy for users from one IdP

You can configure a service provider (SP) authentication policy to enforce authentication requirements for an identity provider (IdP) connection.

Before you begin

This example requires the following components:

- An SP adapter instance deployed, configured, and integrated with the target application.
- An IdP connection to the partner. For more information, seestep 1.
- An IdP connection to the third-party IdP that facilitates the multifactor authentication process. For more information, seestep 2.
- An authentication policy contract to carry user attributes from the partner to the target application. For more information, see step 3.
- An SP authentication policy. For more information, see step 4 and step 7.
- An adapter mapping between the authentication policy contract and the applicable SP adapter instance. For more information, see step 5.
- An SP-initiated single sign-on (SSO) URL. For more information, see step 6.

About this task

In this example, you want to create an IdP connection to Alpha, which passes two attributes in its assertions, SAML SUBJECT and samlEmail, on your PingFederate SP server. You also want to enforce multi-factor authentication (MFA) for users from Alpha through Bravo, a third-party IdP that returns only the SAML SUBJECT attribute and requires a user ID to be passed in from the original source. Both Alpha and Bravo support SAML 2.0 and only the SP-initiated single sign-on (SSO) profile.

Create an SP adapter instance using the Applications # Integration # SP Adapters configuration wizard and complete the last-mile integration with the target application. The SP adapter instance name and ID are Sample and sample, respectively. On System # Server # Protocol Settings # Federation Info, the

base URL for your PingFederate SP server is https://sso.xray.local:9031. There are no other IdP connections besides those required to connect with Alpha and Bravo.

Steps

- 1. Go to Authentication # Integration # IdP Connections. On the IdP Connections window, click **Create Connection.**
 - a. Follow the connection workflow to create a SAML 2.0 IdP connection to Alpha.
 - In this example, Alpha's entity ID is sso.alpha.local.
 - b. On the **SAML Profiles** tab, make sure that the **SP-Initiated SSO** check box is selected.
 - c. On the **Identity Mapping** window, select **No Mapping**.



Tip:

If the partner and your organization agree to support account linking, select **Account Linking**. If the partner and your organization agree to support the IdP-initiated profile, select **Account** Mapping or Account Linking.

Both use cases require the applicable SP adapter instance (Sample) to be mapped into the connection. To get to the Target Session Mapping tab, go to IdP Connection # Browser SSO # **User-Session Creation**. The rest of the steps remain unchanged.

This sample configuration uses neither Account Linking or Account Mapping.

- d. Complete the rest of the connection configuration.
- 2. Repeat step 1 to create a SAML 2.0 IdP connection to Bravo.

In this example, Bravo's entity ID is sso.bravo.local.

3. Go to Authentication # Policies # Policy Contracts. On the Policy Contracts window, click Create **New Contract.**



The purpose of an authentication policy contract is to harness user attributes obtained through one or more authentication sources as the request flows through the applicable authentication policy. It is the medium between the authentication policies and the target applications. In general:

- You map attributes to authentication policy contracts from authentication policies. For more information, see step 4 in this example.
- You map attributes from authentication policy contracts to target applications through adapter mappings. For more information, see step 5 in this example.
- a. On the **Contract Info** tab, enter a name for this authentication policy contract.

In this example, the name of the policy contract is Authenticated.

- b. On the Contract Attributes tab, enter mail under Extend the Contract and click Add.
- c. Complete the rest of the connection configuration.

Result:

When finished, your policy contract has two attributes: subject and mail.

- 4. Go to **Authentication # Policies # Policies**. On the **Policies** window, click **Add Policy** to define a policy to enforce the third-party authentication requirement.
 - a. On the **Policy** window, enter a name and a description for this policy, and select **sso.alpha.local** (**IdP Connection**) as the first policy action from the list.

Result:

An IdP connection has two results, **Fail** and **Success**, as illustrated in the following screen capture.

Each result forms its own policy path that requires further configuration.

b. For the **sso.alpha.local** # **Fail** path, click **Done** as the policy action.

Result:

At runtime, PingFederate terminates the request and returns an error message to the user.

- c. From the **sso.alpha.local** # **Success** list, select **sso.bravo.local** as the policy action.
- d. Below sso.bravo.local, click Options to relay the user ID (SAML_SUBJECT) from sso.alpha.local to sso.bravo.local.
- e. On the **Incoming User ID** dialog window, select **IdP Connection (sso.alpha.local)** from the **Source** list and **SAML SUBJECT** from the **Attribute** list.

To close the **Incoming User ID** dialog window, click **Done**.

f. For the sso.bravo.local # Fail path, select Done as the policy action.

Result:

At runtime, PingFederate terminates the request and returns an error message to the user.

g. For the sso.bravo.local # Success path, select the authentication policy contract created in step 3.

Result:

Your policy should be similar to the following sample.

- h. Below the authentication policy contract, click **Contract Mapping** and follow the **Manage Authentication Policies** # **Authentication Policy Contract Mapping** configuration workflow to configure the fulfillment of the authentication policy contract.
- i. On the **Attribute Sources & User Lookup** tab, click **Add Attribute Source** to configure datastore queries.
- j. On the Contract Fulfillment tab, from the Source list select IdP Connection (sso.alpha.local) and the appropriate attributes from the Value list to fulfill the authentication policy contract attributes.





Tin-

In this configuration, you are mapping attributes to an authentication policy contract from the authentication policy.

Optional: On the **Issuance Criteria** tab, configure conditions to be validated before issuing an authentication policy contract.

On the **Summary** tab, click **Done**.

On the **Policy** window, click **Done**.

On the **Policies** window, click **Save**.

- 5. Go to **Applications** # **Integration** # **Policy Contract Adapter Mappings**. Map the authentication policy contract to the SP adapter instance that you have deployed, configured, and integrated with the target application.
 - a. From the **Source Instance**list, select the authentication policy contract, created in *step 3*.
 - b. From the **Target Instance** list, select the applicable SP adapter instance (**Sample**).
 - c. Click Add Mapping.
 - d. Follow the **Mapping Configuration** workflow to create the mapping.
 - e. On the **Attribute Sources & User Lookup** tab, click **Add Attribute Source** to configure datastore queries.
 - f. On the Adapter Contract Fulfillment window, select Authentication Policy Contract from the Source list and the appropriate contract attributes from the Value list to fulfill the SP adapter contract.



Tip:

In this configuration, you are mapping attribute values from the authentication policy contract to the target application through the applicable SP adapter instance.

Optional: On the **Default Target URL** tab, specify a default target URL for this mapping configuration.

Optional: On the **Issuance Criteria** tab, configure conditions to be validated before issuing an SP adapter contract.

Click Done.

6. Configure an SP-initiated SSO URL in your target application by combining the base URL of your PingFederate SP server, https://sso.xray.local:9031, the PingFederate's application SSO endpoint, /sp/startSSO.ping, and the SpSessionAuthnAdapterId parameter with the adapter ID of the applicable SP adapter instance as the parameter value.

For example: https://sso.xray.local:9031/sp/startSSO.ping?SpSessionAuthnAdapterId=sample

If you have not defined a default URL for the adapter mapping, configured in *step 5*, the IdP connection, or the PingFederate SP server, you must also configure your target application to include the **TargetResource** parameter in its SP-initiated SSO requests.



Important:

When using the parameters <code>TargetResource</code> or <code>TARGET</code> with their own query parameters included, the parameter value must be URL-encoded. Any other parameters that contain restricted characters, such as many SAML URNs, also must be URL-encoded. For information about URL encoding, see third party resources such as HTML URL-encoding Reference. Parameters are case-sensitive.

7. When you are ready to test your new use case, go to **Authentication # Policies # Policies**. On the **Policies** window, select the **SP Authentication Policies** check box to enable policies for SP-initiated Browser SSO requests received by at the /sp/startsso.ping endpoint, and click **Save**.

You can configure service provider (SP) authentication policies to handle different authentication requirements for multiple identity provider (IdP) connections.

About this task

Assume you configure the following use cases in an earlier version of PingFederate:

Two SP adapter instances on Applications # Integration # SP Adapters.

Instance Name	Instance ID	Extended Contract
Sample	sample	subject and email
Sample Delta	sampleDelta	subject and email

Three entries on Applications # Integration # Target URL Mapping.

URL	Target Session
https://sso.xray.local:9031/SpSample/MainPage?app=Alpha&*	Sample
https://sso.xray.local:9031/SpSample/MainPage?app=Charlie&*	Sample
https://sso.xray.local:9031/SpSample/MainPage?app=Delta&*	Sample Delta

Three IdP connections to your partners.

Partner (Federation ID)	Identity Mapping	Attribute Contract	Target Session Mapping
(* ************************************			SP adapter instance name
			(SP adapter instance ID)
Alpha	Account Mapping	SAML_SUBJECT and	Sample
(sso.alpha.local)		samlEmail	(sample)
Charlie	Account Mapping	SAML_SUBJECT and	Sample
(sso.charlie.local)		samlEmail	(sample)
Delta	Account Mapping	SAML_SUBJECT and	Sample Delta
(sso.delta.local)		samlEmail	(sampleDelta)

In this example, all partners support SAML 2.0 and only the SP-initiated single sign-on (SSO) profile.

SP-initiated SSO URLs for users from Alpha, Charlie, and Delta.

Partner	SSO URL
Alpha	https://sso.xray.local:9031/sp/startSSO.ping? PartnerIdpId=sso.alpha.local&TargetResource=https%3A%2F%2Fsso.xray.local %3A9031%2FSpSample%2FMainPage%3Fapp%3DAlph%26t%3Daa
Charlie	https://sso.xray.local:9031/sp/startSSO.ping? PartnerIdpId=sso.charlie.local&TargetResource=https%3A%2F%2Fsso.xray.local%3A9031%2FSpSample%2FMainPage%3Fapp%3DCharlie%26t%3Dc
Delta	https://sso.xray.local:9031/sp/startSSO.ping? PartnerIdpId=sso.delta.local&TargetResource=https%3A%2F%2Fsso.xray.local %3A9031%2FSpSample%2FMainPage%3Fapp%3DDelta%26t%3Dd

- Create new IdP connections to three new partners: Echo, Foxtrot and Golf.
- Enforce multi-factor authentication (MFA) for users from Alpha, Charlie, Echo, and Golf through Bravo.

Bravo requires a user ID to be passed in from the original source and returns only the user ID when the users fulfill the multi-factor authentication (MFA) requirement.

The new required components are:

- Two additional SP adapter instances. For more information, seestep 1:
 - Sample Echo to integrate with Echo's target application.
 - Sample Golf to integrate with Golf's target application.
- Four new IdP connections. For more information, see step 2, step 3, and step 4:

Partner (Federation ID)	Identity Mapping	Attribute Contract	Target Session Mapping	
(i duoi diiidii 12)			SP adapter instance name	
			(SP adapter instance ID)	
Bravo	No Mapping	SAML_SUBJECT and no	N/A	
(sso.bravo.local)		other attributes		
Echo	No Mapping	SAML_SUBJECT and	N/A	
(sso.echo.local)		samlEmail		
Foxtrot	Account Mapping	SAML_SUBJECT and	Sample	
(sso.foxtrot.local)		samlEmail	(sample)	
Golf	No Mapping	SAML_SUBJECT and	N/A	
(sso.golf.local)		samlEmail		

In this example, all partners support SAML 2.0 and only the SP-initiated SSO profile.

- Three authentication policy contracts. For more information, see step 5:
 - An authentication policy contract, Authenticated, to carry user attributes from Alpha and Charlie to their respective target applications.
 - Two other authentication policy contracts, Echo authenticated and Golf authenticated, to carry
 user attributes from Echo and Golf to their target applications.
- An instance of the HTTP Request Parameter Authentication Selector, Partnerldpld, to determine if a request is meant for Alpha or Charlie, because Alpha's and Charlie's target applications share an SP adapter instance. For more information, see step 6.
- Three SP authentication policies to enforce the multifactor authentication requirement. For more information, see step 7, step 8, and step 12.
- Three adapter mappings between the authentication policy contracts and the applicable SP adapter instances. For more information, see step 9:
 - Map from Authenticated to Sample.
 - Map from Echo authenticated to Sample Echo.
 - Map from Golf authenticated to Sample Golf
- Three additional target URL mappings between the applications requested by users from Echo,
 Foxtrot, and Golf to their respective SP adapter instances. For more information, see step 10:
- SSO URLs for all partners. For more information, see step 11.

Follow these steps to fulfill the new requirements:

Steps

1. Go to Applications # Integration # SP Adapters. On the SP Adapters window, create two new SP adapter instances, as shown in the following table.

Instance Name	Instance ID	Extended Contract
Sample Echo	sampleEcho	subject and email
Sample Golf	sampleGolf	subject and email

- Create an IdP connection to Bravo.
 - a. On the IdP Connections window (Authentication # Integration # IdP Connections) create a SAML 2.0 IdP connection.
 - b. On IdP Connection # Browser SSO # SAML Profiles, select the SP-Initiated SSO check box.
 - c. On IdP Connection # Browser SSO # User-Session Creation # Identity Mapping, select No Mapping.
 - d. Complete the rest of the connection configuration.
- 3. Create IdP connections to Echo and Golf.
 - a. Go to Authentication # Integration # IdP Connections. On the IdP Connections window, create a SAML 2.0 IdP connection to Echo, and then to Golf.
 - b. On IdP Connection # Browser SSO # SAML Profiles, select the SP-Initiated SSO check box.
 - c. On IdP Connection # Browser SSO # User-Session Creation # Identity Mapping, select No Mapping.



If the partner and your organization agree to support account linking, select **Account Linking**. When the **Account Linking** option is selected, you must map the applicable SP adapter instance, Sample Echo for Echo or Sample Golf for Golf, to the connection on IdP Connection # Browser SSO # User-Session Creation # Target Session Mapping. The rest of the steps remain unchanged.

This example does not use account linking.

- d. On IdP Connection # Browser SSO # User-Session Creation # Attribute Contract, enter samlEmail under Extend the Contract.
- e. Complete the rest of the connection configuration.
- f. Repeat these steps to create a SAML 2.0 IdP connection to Golf.
- 4. Create an IdP connection to Foxtrot.
 - a. On the IdP Connections window, create a SAML 2.0 IdP connection.
 - b. On IdP Connection # Browser SSO # SAML Profiles, select the SP-Initiated SSO check box.
 - c. On the User-Session Creation tab, click Configure User-Session Creation.
 - d. On the Identity Mapping tab, select Account Mapping.
 - e. On the Attribute Contract tab, enter samlEmail under Extend the Contract.
 - f. On the Target Session Mapping tab, click Map New Adapter Instance and follow the Adapter Mapping & User Lookup configuration workflow to map the attributes from the assertion to the SP adapter instance Sample.

Example:

Adapter Contract	Source	Value
email	Assertion	samlEmail

Adapter Contract	Source	Value
subject	Assertion	SAML_SUBJECT

- g. Complete the rest of the connection configuration.
- 5. Create three authentication policy contracts, one for Alpha and Charlie, one for Echo, and one for Golf.



The purpose of an authentication policy contract is to harness user attributes obtained through one or more authentication sources as the request flows through the applicable authentication policy. It is the medium between the authentication policies and the target applications. In general:

- You map attributes to authentication policy contracts from authentication policies. For more information, seestep 7 in this example.
- You map attributes from authentication policy contracts to target applications through adapter mappings. For more information, see step 9 in this example.
- a. Go to Authentication # Policies # Policy Contracts. On the Policy Contracts window, click Create New Contract to create an authentication contract for users from Alpha and Charlie, for users from Echo, and then for users from Golf.
- b. On the Contract Info tab, enter a name for this authentication policy contract.

In this example, the names are Authenticated, Echo authenticated, and Golf authenticated.

- c. On the Contract Attributes tab, extend the authentication policy contract with an attribute for user's email address, such as mail.
- d. Complete the rest of the connection configuration.
- e. Repeat these steps to create an authentication policy contract for users from Echo, and then for users from Golf.

Result:

Your policy contracts should be similar to the following samples.



6.



Note:

If multiple target applications share the same SP adapter instance, you must use a selector to evaluate the SP-initiated requests such that the policy engine can route the requests to the policy paths that are meant for the respective IdPs. This example uses an instance of the HTTP Request Parameter

Create an instance of the HTTP Request Parameter Authentication Selector.

- a. Go to Authentication # Policies # Selectors. On the Selectors window, click Create New Instance.
- b. On the **Type** tab, select **HTTP Request Parameter Authentication Selector** from the **Type** list and provide a name and ID for the selector instance.
 - In this example, the name and ID are both PartnerIdpId.
- c. On the Authentication Selector tab, enter PartnerIdpId in the HTTP Request Parameter Name field.
- d. On the Selector Result Values tab, enter sso.alpha.local and sso.charlie.local as the result values.



Note:

In general, for the IdPs that you want to enforce additional authentication requirements through one or more SP authentication policies and whose target applications share an SP adapter instance, you must enter their federation IDs here.

e. Complete the rest of the configuration.

Result:

Your selector instance should be similar to the following sample.



7. Go to Authentication # Policies # Policies. On the Policies window, click Add Policy to define a policy to enforce the third-party authentication requirement for users from Echo, and then for users from Golf.



Tip:

If you need more information about each sub step, see step 4 in Configuring an SP authentication policy for users from one IdP on page 256.

- a. On the Policy window, enter a name and a description for this policy, and select sso.echo.local (IdP Connection) as the first policy action from the list.
- b. For the **sso.echo.local** # **Fail** path, select **Done** as the policy action.
- c. For the sso.echo.local # Success path, select sso.bravo.local (IdP Connection) as the policy action.
- d. Click Options, underneath sso.bravo.local (IdP Connection), to relay the user ID (SAML SUBJECT) from sso.echo.local to sso.bravo.local.
- e. For the **sso.bravo.local** # **Fail** path, select **Done** as the policy action.
- f. For the sso.bravo.local # Success path, select the policy contract Echo authenticated as the policy action, and then click Contract Mapping, underneath the policy contract, to configure the fulfillment of the policy contract.



Essentially, you are mapping attributes to an authentication policy contract from the authentication policy.

g. Repeat these steps to define a new authentication policy to enforce third-party authentication for users from Golf.

Result:

Your policies should be similar to the following samples.

8. On the **Policies** tab, click **Add Policy** to define a new authentication policy to enforce third-party authentication for users from Alpha and Charlie.



Note:

If multiple target applications share the same SP adapter instance, you must use a selector to evaluate the SP-initiated Browser SSO requests, such that the policy engine can route the requests to the

- a. On the **Policy** window, enter a name and a description for this policy, and select the instance of the HTTP Request Parameter Authentication Selector as the first policy action from the list.
- b. For the **sso.alpha.local** path, repeat step 7 to configure the authentication requirement for the sso.alpha.local IdP connection.
- c. For the **sso.charlie.local** path, repeat step 7 to configure the authentication requirement for the sso.charlie.local IdP connection.
- d. For the **sso.foxtrot.local** path, repeat step 7 to configure the authentication requirement for the sso.foxtrot.local IdP connection.

Result:

Your policy should be similar to the following sample.



Note:

This window capture collapses the sso.charlie.local and sso.foxtrot.local policy paths for documentation presentation purposes.

When you go back to the **Policies** tab, your policies should be similar to the following samples.

e. Click Save.

- 9. Create adapter mappings.
 - a. Go to Applications # Integration # Policy Contract Adapter Mappings. On the Policy Contracts window, create three adapter mappings to map from the authentication policy contracts, created in *step 5*, to the applicable SP adapter instances.
 - Map from Authenticated to Sample.
 - Map from Echo authenticated to Sample Echo.
 - Map from Golf authenticated to Sample Golf.



If you need more information about each sub step, see *step 5* in *Configuring an SP authentication policy for users from one IdP* on page 256.

Your mappings should be similar to the following samples.



Each adapter mapping should be similar to the following configuration.





You are essentially mapping attribute values from the authentication policy contracts to target applications through the applicable SP adapter instances.

- 10. Create target URL mappings.
 - a. Go to **Applications # Integration # Target URL Mapping**. On the **Target URL Mapping** window, add the following mappings.

URL	Target Session
https://sso.xray.local:9031/SpSample/MainPage/?app=Echo&*	Sample Echo
https://sso.xray.local:9031/SpSample/MainPage/?app=Foxtrot&*	Sample
https://sso.xray.local:9031/SpSample/MainPage/?app=Golf&*	Sample Golf

Result:

Your target URL mappings should be similar to the following samples.

11. Configure the following SSO URLs.

Partner	SSO URL
Alpha	https://sso.xray.local:9031/sp/startSSO.ping? PartnerIdpId=sso.alpha.local&TargetResource=https%3A%2F %2Fsso.xray.local%3A9031%2FSpSample%2FMainPage%3Fapp%3DAlpha %26t%3Da
	The SSO URL has not changed.
	Based on the current configuration, because the target applications for Alpha and Charlie share the same SP adapter instance, the PartnerIdpId query parameter is required for the configured policy to route the request to the corresponding IdP connection.
Charlie	https://sso.xray.local:9031/sp/startSSO.ping? PartnerIdpId=sso.charlie.local&TargetResource=https%3A%2F %2Fsso.xray.local%3A9031%2FSpSample%2FMainPage%3Fapp%3DCharlie %26t%3Dc
	The SSO URL has not changed.
	Based on the current configuration, because the target applications for Alpha and Charlie share the same SP adapter instance, the PartnerIdpId query parameter is required for the configured policy to route the request to the corresponding IdP connection.
Delta	https://sso.xray.local:9031/sp/startSSO.ping? PartnerIdpId=sso.delta.local&TargetResource=https%3A%2F %2Fsso.xray.local%3A9031%2FSpSample%2FMainPage%3Fapp%3DDelta %26t%3Dd
	The SSO URL has not changed.
	Optionally, based on the current configuration, you can remove the <code>PartnerIdpId</code> query parameter because it is not required. You can also replace the <code>TargetResource</code> query parameter and its value with the <code>SpSessionAuthnAdapterId</code> query parameter and the applicable SP adapter instance ID, <code>SpSessionAuthnAdapterId=sampleDelta</code> , if you have configured a default target URL in the IdP connection to Delta, or a default SP SSO URL for all IdP connections.
Echo	https://sso.xray.local:9031/sp/startSSO.ping? SpSessionAuthnAdapterId=sampleEcho&TargetResource=https%3A%2F %2Fsso.xray.local%3A9031%2FSpSample%2FMainPage%3Fapp%3DEcho %26t%3De
	This is a new SSO URL.
	Optionally, based on the current configuration, you can remove the <code>TargetResource</code> query parameter and its value if you have configured a default target URL in the IdP connection to Echo, or a default SP SSO URL for all IdP connections.



Important:

When using the parameters TargetResource or TARGET with their own query parameters included, the parameter value must be URL-encoded. Any other parameters that contain restricted characters, such as many SAML URNs, also must be URL-encoded. For information about URL encoding, see third party resources such as HTML URL-encoding Reference. Parameters are case-sensitive.

12. To test your new use case, go to Authentication # Policies # Policies. On the Policies window, select the SP Authentication Policies check box to enable policies for SP-initiated Browser SSO requests received by at the /sp/startSSO.ping endpoint, and click Save.

Configuring SP authentication policies for internal users

The /pf/adapter2adapter.ping endpoint initiates direct IdP-to-SP adapter mapping, mostly intended for the internal users to access resources without maintaining a service provider (SP) and an identity provider (IdP) connection on the same server.

About this task

To prevent users from circumventing the SP authentication policies, this endpoint becomes inactive when SP authentication policies are enabled but IdP authentication policies are disabled. Administrators can configure SP authentication policies for the internal users to re-enable access to protected resources.

Suppose you have configured the following use cases in PingFederate 8.0.

For users from Hotel, an IdP:

- An SP adapter instance:
 - Name: Sample Hotel
 - ID: sampleHotel

- A SAML 2.0 IdP connection:
 - Partner: Hotel
 - Federation ID: sso.hotel.local
 - SAML Profile: SP-initiated single sign-on (SSO) only
 - Identity mapping method: Account mapping
 - Default target URL: https://sso.xray.local:9031/SpSample/MainPage/?app=Hotel&t=h
 - SSO URL: https://sso.xray.local:9031/sp/startSSO.ping?PartnerIdpId=sso.hotel.local

For internal users:

- An IdP HTML Form Adapter instance, HTML Form, validating credentials through a Password Credential Validator (PCV) instance against your user directory
- An adapter-to-adapter mapping:
 - Source: HTML FormTarget: Sample Hotel
 - Default target URL: https://sso.xray.local:9031/SpSample/MainPage/?app=Internal&t=i
 - SSO URL: https://sso.xray.local:9031/pf/adapter2adapter.ping?
 SpSessionAuthnAdapterId=sampleHotel

After upgrading to PingFederate 10.1, if you want to enforce multi-factor authentication for users from Hotel through Bravo, you can create an IdP connection to Bravo and the following authentication policy.

Because the authentication policy ends with a policy contract **Hotel authenticated**, you must create an adapter mapping, from the policy contract, to **Sample Hotel**, the SP adapter instance integrated with the target application. You also need to update the SSO URL for users from Hotel to https://sso.xray.local:9031/sp/startSSO.ping?SpSessionAuthnAdapterId=sampleHotel.

When you select the **SP** Authentication Policies check box without selecting the **IdP** Authentication Policies check box, the /pf/adapter2adapter.ping endpoint is disabled to prevent malicious Hotel's users, with specific knowledge of PingFederate endpoints, your PingFederate configuration, and functional credentials, from trying to access the target application through the SSO URL intended for your internal users. By doing so, you circumvent the SP authentication policy that is meant for them.

To re-enable access to the application for the internal users, you need the following new components:

- An additional SP adapter instance, Sample Internal to integrate with the target application. See step
- An authentication policy contract, **Internal authenticated**, to carry attributes from internal users to the target applications. See *step 2*.
- An instance of the CIDR Authentication Selector to be deployed in the authentication policy to reject users from external networks to access protected resources using your HTML Form Adapter. Alternatively, deploy an instance of the PingID Adapter in the authentication policy to enforce multifactor authentication for users who authenticated successfully using the HTML Form Adapter. See step 3
- An authentication policy for the internal users. See step 4.
- An adapter mapping to map from the authentication policy contracts Internal authenticated to the SP adapter instance Sample Internal. See <u>step 5</u>.
- A new SSO URL for the internal users. See step 6.

Follow these steps to fulfill the new requirements.

Steps

1. Go to Applications # Integration # SP Adapters.

For each new instance, include the name, ID, and extended contract.

Instance Name	Instance ID	Extended Contract
Sample Internal	sampleInternal	subject and email

- 3. Go to Authentication # Policies # Policy Contracts.
- 4. Click Create New Contract to create an authentication policy contract.



In this example, the name of the policy contract is Internal authenticated.

The purpose of an authentication policy contract is to harness user attributes obtained through one or more authentication sources as the request flows through the applicable authentication policy. It is the medium between the authentication policies and the target applications. Generally speaking:

- You map attributes to authentication policy contracts from authentication policies. step 4 in this example.
- You map attributes from authentication policy contracts to target applications through adapter mappings, step 5 in this example.
- 5. Go to Authentication # Policies # Selectors.
- 6. Click Create New Instance to create an instance of the CIDR Authentication Selector with one or more network ranges that correspond to your internal users.



Note:

In this example, the name and ID are both Internal users.

The purpose of the CIDR Authentication Selector is for the policy engine to reject users from external networks to access protected resources using your HTML Form Adapter. You can also deploy the PingID Adapter to enforce multifactor authentication for users authenticated through the HTML Form Adapter. If users fail to fulfill the PingID multifactor authentication requirement, the policy engine rejects their requests, thus providing another layer of protection against unauthorized access from malicious users.

7. Go to Authentication # Policies # Policies to configure your policies as follows.

With PingID Adapter

With CIDR Authentication Selector

- 8. Go to Applications # Integration # Policy Contract Adapter Mappings.
- 9. Click Create New Instance and create a mapping from the new authentication policy contract, Internal authenticated., to the new SP adapter instance, sampleInternal.
- 10. On the Default Target URL window, enter https://sso.xray.local:9031/SpSample/ MainPage/?app=Internal&t=i.
- 11. Update the SSO URL for your internal users to https://sso.xray.local:9031/sp/startSSO.ping? SpSessionAuthnAdapterId=sampleInternal.

Policy contracts

Authentication policy contracts, formerly known as connection mapping contracts, give PingFederate administrators a variety of capabilities.

Authentication policy contracts provide PingFederate administrators the following benefits:

- The capability to build an attribute contract with attribute values from multiple authentication sources or datastore queries through an authentication policy.
- The flexibility to map only the policy contract to a connection. Administrators do not have to map into the connection the authentication sources in the policy leading up to the contract. For example, administrators can experiment with various IdP adapter instances without the burden of adding and removing them to and from the connection.
- The potential to reuse authentication policies that use the same policy contract in multiple service provider (SP) connections, identity provider (IdP) connections, and OAuth use cases, using the OAuth Authorization Code or Implicit grant types.

Authentication policy contracts are also the media to carry user attributes from IdPs to SPs when PingFederate is deployed as a federation hub. For more information, see *Federation hub use cases* on page 99.

Managing policy contracts

Manage authentication policy contracts from the Authentication tab.

About this task

You manage authentication policy contracts in the **Authentication # Policies # Policy Contracts** window.

Steps

- 1. Click **Create New Contract** to create a new authentication policy contract.
- 2. Edit an existing authentication policy contract by clicking its name.
- 3. Check the policy contract usage.
- 4. Optional: On the **Policy Contracts** window, click **Delete** to remove an authentication policy contract, if the authentication policy contract is not in use, or cancel the removal request.

Editing contract information

Edit contract information on the **Contract Info** window.

Steps

- 1. Go to Authentication # Policies # Policy Contracts.
- 2. Click Create New Contract or select from an available contract.
- 3. On the **Contract Info** window, enter or modify the contract name.
- 4. Click Done.

Defining contract attributes

Every authentication policy contract comes with a subject attribute. You can extend the contract with additional attributes as needed.

About this task

To manage the user attributes in the authentication policy contract, go to **Authentication # Policies # Policy Contracts**.

If needed, enter an attribute under Extend the Contract, and then click Add.



Note:

Attribute names are case-sensitive and must suit the needs of your partners. Repeat to add more attributes as needed.

- Click **Edit**, **Update**, or **Cancel** to modify an existing attribute or undo changes.
- Click **Delete** to remove an existing attribute.

Reviewing the policy contract

Make changes to any policy contract you have added.

Steps

- To keep your changes, click Save.
- To amend your configuration, click the name of the corresponding tab and then follow the configuration wizard to complete the task.
- To discard your changes, click Cancel.

Adapter Mappings

Configuring adapter mappings allows administrators to map attributes from an authentication policy contract directly to a service provider (SP) adapter instance.

This allows the administrators to chain multiple authentication sources in an SP authentication policy, to build an authentication policy contract using attributes from authentication sources in the policy path, and to apply the authentication policy contract to the target application.

Configuring authentication policy adapter mappings

Authentication policy adapter mappings allow administrators to map attributes from an authentication policy contract directly to a service provider (SP) adapter instance.

Steps

- 1. Go to Applications # Integration # Policy Contract Adapter Mappings.
- 2. From the **Source Instance** list, select the applicable authentication policy contract from the .
- 3. From the **Target Instance** list, select the SP adapter instance integrated with your target application.
- 4. Click Add Mapping.

- 5. Follow the Applications # Integration # Adapter-to-Adapter Mappings wizard to create the mapping.
 - a. Optional: On the Attribute Sources & User Lookup tab, click Add Attribute Source to configure datastore queries to fulfill the SP adapter contract.
 - Queries are executed in the order they are displayed on the Attribute Sources & User Lookup tab. Use the up and down arrows as needed to adjust the order.
 - If a required attribute cannot be fulfilled, such as the user identifier of an adapter, the request fails. For more information, see *Fulfillment by datastore gueries* on page 361.
 - b. On the Adapter Contract Fulfillment tab, select a source and an attribute to fulfill the SP adapter contract.



Select Authentication Policy Contract from the Source list to map directly from the policy contract to the SP adapter contract or another choice to fulfill the SP adapter contract through datastore queries, dynamic texts, or results from OGNL expression

- c. Optional: On the **Default Target URL** tab, specify a default target URL for this mapping configuration.
- d. Optional: On the Issuance Criteria tab, configure conditions to be validated before issuing an SP adapter contract. For more information, see Define issuance criteria for adapter mapping.
- e. On the Adapter-to-Adapter Mapping Summary tab, review the configuration and modify as needed. When complete, click Done.
- 6. On the Adapter-to-Adapter Mappings window, click Save.

Defining issuance criteria for adapter mapping

You can define issuance criteria for adapter mappings in PingFederate to conditionally approve or reject requests based on individual attributes. Satisfy criteria in order for PingFederate to move a request to the next phase.

About this task

On the **Issuance Criteria** tab, define the criteria to satisfy in order for PingFederate to further process a request. Use this token authorization feature to conditionally approve or reject requests based on individual attributes.

Begin this optional configuration by choosing the source that contains the attribute to verify. Some sources are common to almost all use cases, such as Mapped Attributes. Other sources depend on the type of configuration, such as **JDBC**. Irrelevant sources are automatically hidden. Once you select a source, choose the attribute to verify. Depending on the selected source, the available attributes or properties vary. Specify the comparison condition and the desired value to compare to.

You can define multiple criteria, which must all be satisfied in order for PingFederate to move a request to the next phase. A criterion is satisfied when the runtime value of the selected attribute matches or does not match the specified value, depending on the chosen comparison method. The multi-value contains ... or multi-value does not contain ... comparison methods are intended for attributes that can contain multiple values. Such a criterion is considered satisfied if one of the multiple values match or does not match the specified value. Values are compared verbatim. If you require complex evaluations, including conditional criteria or partial matching, define them using attribute mapping expressions.



Note:

All criteria defined must be satisfied, or evaluated as true, for a request to move forward, regardless of how the criteria were defined. As soon as one criterion fails, PingFederate rejects the request and returns an error message.

1. From the **Source** list, select the attribute's source.

Depending on the selection, the Attribute Name list populates with associated attributes. See the following table for more information.

Source	Description
Authentication Policy Contract	Select to evaluate attributes from the authentication policy contract.
Context	Select to evaluate properties returned from the context of the transaction at runtime.
	Note:
	The HTTP Request context value is retrieved as a Java object rather than text. For this reason, attribute mapping expressions are more appropriate to evaluate and return values.
JDBC, LDAP, or other types of datastore (if configured)	Select to evaluate attributes returned from a data source.
Mapped Attributes	Select to evaluate the mapped attributes.

- 2. From the Attribute Name list, select the attribute to be evaluated.
- 3. From the **Condition** list, select the comparison method.

Available methods:

- equal to
- equal to (case insensitive)
- equal to DN
- not equal to
- not equal to (case insensitive)
- not equal to DN
- multi-value contains
- multi-value contains (case insensitive)
- multi-value contains DN
- multi-value does not contain
- multi-value does not contain (case insensitive)
- multi-value does not contain DN



The first six conditions are intended for single-value attributes. Use one of the multi-value ... conditions for PingFederate to validate whether one of the attribute values matches the specified value. When an attribute has multiple values, using a single-value condition causes the criteria to fail.



Values are compared verbatim. If you require complex evaluations, including conditional criteria or partial matching, define them using attribute mapping expressions. For more information, see Attribute mapping expressions on page 214.

5. In the **Error Result** field, enter a custom error message.

To use localized descriptions, enter a unique alias in the Error Result field, such as someIssuanceCriterionFailed. Insert the same alias with the desired localized text in the applicable language resource files, located in the <pf install>/pingfederate/server/ default/conf/language-packs directory.

If not defined, PingFederate returns ACCESS DENIED when the criterion fails at runtime.

- 6. Click Add.
- 7. Optional: Repeat to add more criteria.
- 8. If you require complex evaluations, including conditional criteria or partial matching, define them using attribute mapping expressions. For more information, see Attribute mapping expressions on page 214.
 - a. Click Show Advanced Criteria.
 - b. In the **Expression** field, enter the required expressions.
 - c. Optional: In the **Error Result** field, enter an error code or message.



Note:

If the expressions resolve to a string value instead of true or false, the returned value overrides the Error Result field value.

- d. Click Add.
- e. Optional: Click **Test**, enter values in the applicable fields, and verify the results.
- f. Optional: Repeat to add multiple criteria using attribute mapping expressions.

Sessions

This topic describes the differences between application and authentication sessions.

Application sessions

Application sessions apply to PingFederate applications hosted on its user-facing endpoints, such as the profile management page and the grant management endpoints. When the inactivity threshold or the maximum lifetime is reached, PingFederate redirects previously authenticated users back to the authentication sources, identity provider (IdP) adapter instances or IdP connections, subject to the configuration of authentication sessions.

Authentication sessions

Authentication sessions control when PingFederate redirects previously authenticated users back to the authentication sources on subsequent requests for browser-based single sign-on (SSO) or PingFederate applications.

Authentication sessions typically wrap an adapter so that PingFederate creates the session when user authentication has succeeded. PingFederate invokes the adapter's authentication logic again only when the session reaches its limits. However, depending on the implementation, an adapter can be aware of an authentication session that wraps it and override this logic. In particular, PingFederate creates authentication sessions configured for an Identifier First Adapter instance only when the complete single sign-on (SSO) transaction has succeeded. This lets the adapter prompt the user for a different user identifier when a chained adapter authentication fails because, for example, there's a typo in the user identifier.

Session storage options

When authentication sessions are enabled, PingFederate maintains session data in memory.

PingFederate also supports maintaining session data both in memory and on an external storage. This optional capability allows administrators to support use cases where a longer session duration or a greater resilience against restarts of PingFederate and browsers is desired. The retrieval and update operations are optimized to provide a fast and seamless user experience. For instance, a retrieval from the external storage is only required when an authentication session is not found in memory.



Persistent authentication sessions require an external storage. For more information, see Defining a datastore for persistent authentication sessions on page 884.

Inactivity (idle) timeout and maximum lifetime

When authentication sessions are enabled, an authenticated user is not sent back to the authentication system as long as the user makes another request within the idle timeout window, 60 minutes by default. If the user makes another request within the idle timeout window, the authentication session is extended by the idle timeout value, another 60 minutes by default. For externally stored authentication sessions, this operation is optimized to only send updates to the external storage when the remaining idle timeout window is less than 75%.

An authentication session can be repeatedly extended by multiple requests and remains valid until the maximum timeout value is reached, in which case the user will be redirected back to the authentication system.



The authentication system might or might not challenge the user to complete an authentication process based on its own session management policy or processing logic.

Configuration options

Administrators can enable authentication sessions for all authentication sources, with or without making the authentication sessions persistent, and with or without specifying overrides for selected authentication sources.

Alternatively, administrators can enable authentication sessions for a few selected authentication sources, optionally with their own sets of overrides. The override options include:

- Disable or enable authentication sessions
- Make authentication sessions persistent
- Override the idle timeout, the maximum timeout, or both, in minutes, hours, or days
- Enforce authentication requirement based on authentication context

Because sessions are tracked with their respective authentication context, administrators can optionally configure PingFederate to compare the requested authentication context found in the authentication request against the authentication context found in the session. If the values do not match, PingFederate redirects the user back to the authentication system.

Tracking options for logout

Administrators can optionally configure additional tracking options for logout to control whether PingFederate should leverage the single logout (SLO) application endpoints to terminate adapter sessions, add sessions to the session revocation list as users sign out, or do both. Publish revoked sessions to provide a secure SLO experience with PingAccess deployments.

Configuring tracking options for logout

You can configure PingFederate to track adapter sessions for logout.

About this task

An adapter session is a logout entry that, if tracked, ensures a logout request is sent to the adapter during single logout (SLO). Then the adapter can remove any session data that it is tracking for the user.

Steps

- 1. Go to Authentication # Policies # Sessions.
- 2. Optional: Enable SLO for all adapter instances on a per-user basis by selecting the Track Adapter Sessions for Logout check box.

When this check box is selected, an adapter session is tracked whenever an adapter is invoked during single sign-on (SSO). When this check box is not selected, the tracking of the adapter session depends on other factors, such as whether SLO is enabled on the partner connection involved in the SSO. This check box is not selected by default.

3. Optional: Add the associated sessions to the revocation list on logout by selecting the Track Revoked Sessions on Logout check box.

When selected, PingFederate always adds the associated sessions to the session revocation list as users sign off, even if an error occurs to the logout requests. This allows other systems, such as PingAccess, to guery the validity of a given session at the Session Revocation API endpoint, /pf-ws/ rest/sessionMgmt/revokedSris. This check box is selected by default for new installations.



Note:

If your use cases involve OAuth requests, consider enabling the Check session revocation status option in the applicable Access Token Management instances so that the token validation process takes into account whether a session has been added to the revocation list. For more information, see Managing session validation settings on page 557.

4. Optional: Change the number of minutes until the revoked sessions are removed from the revocation list for optimal performance by changing the value in the Session Revocation Lifetime field. You can enter an integer between 1 and 43200. The default value is 490 minutes.



Important:

The Session Revocation Lifetime value should match or exceed the idle timeout value, or the maximum session lifetime value, of the authentication sources and the relying parties. For example, the default value of 490 minutes exceeds the global Max Timeout value for authentication sessions by 10 minutes to allow for clock skew among servers.

5. Click Save.

You can configure and override the default timeout limits for application sessions in the PingFederate administrative console.

Steps

- 1. Go to Authentication # Policies # Sessions.
- 2. Optional: On the Sessions window, override the default timeout values under Application Sessions.

Field	Description
Idle Timeout (Minute)	Modify the default inactivity timeout value in the Idle Timeout (Minute) field.
	You can enter an integer between 1 and 1576800 , representing a range of one minute to 1,095 days. You can also empty the value to indicate that the inactivity timeout value should match the maximum lifetime.
	The default value is 60 minutes.
Max Timeout (Minutes)	Modify the default maximum lifetime of an authentication session in the Max Timeout (Minutes) field.
	You can enter an integer between 1 and 1576800, representing a range of one minute to 1,095 days. You can also empty the value to indicate that the authentication sessions do not expire until they are removed from the system.
	The value of the Max Timeout (Minutes) field cannot be less than that of the Idle Timeout (Minute) field.
	The default value is 480 minutes, which is 8 hours.

3. To keep your configuration, click Save.

Configuring authentication sessions

You can configure and override the default timeout limits for authentication sessions in the PingFederate administrative console.

Steps

- 1. Go to Authentication # Policies # Sessions.
- 2. Optional: On the **Sessions** window, configure the global policy and timeout settings under **Authentication Sessions.**
 - a. Select the Enable Sessions for All Authentication Sources check box if PingFederate should track authentication sessions for all authentication sources. Clear this check box if you prefer to enable authentication sessions for only a few authentication sources or disable authentication sessions altogether. This check box is not selected by default.



Note:

For any HTML Form Adapter instance that has been configured to allow users to indicate whether their device is shared or private, if a user signs on without selecting the This is my device check

box on the login form, PingFederate removes authentication session information, if found, and does not store authentication sessions for the user.

b. If your use cases require longer sessions or greater resilience against restarts of PingFederate and browsers, select the Make Authentication Sessions Persistent check box.



Note:

Selecting the check box causes the PF.PERSISTENT cookie to be set in the user's browser. By default, this cookie persists across browser restarts. To allow for very long sessions, the expiration period for the cookie defaults to 94608000 seconds, or 3 years. You can change this period in the cookie-max-age setting in the file persistent-session-cookie-config.xml. If you prefer to have the PF.PERSISTENT cookie cleared on browser exit, set cookie-max-age to -1. Regardless of the cookie's expiration period, PingFederate always enforces the configured session timeouts. However a user might lose their session earlier if the PF.PERSISTENT cookie expires or is removed by the browser.

Persistent authentication sessions require an external storage.



Note:

As of version 9.3, PingFederate alleviates DoS attacks by protecting the persistent session process. It does this by treating repeated persistent cookies that do not have a PF cookie as a replay if repeated in a specified time. This time is set to 300 seconds by default, and you can change it by modifying EmptySessionReplayRetentionsSecs in the <pf-install>/server/default/data/config-store/ org.sourceid.sam120.service.session.StoredSessionServiceImpl.xml file.

For example:

- If a request arrives with a PF.PERSISTENT cookie and without a PF cookie, PingFederate starts counting the time set in EmptySessionReplayRetentionsSecs.
- If another request arrives with the same PF.PERSISTENT cookie and without a PF cookie within the time specified in the configuration file, PingFederate treats it as a replayed request and does not perform a database lookup.

You can disable this behavior by setting EmptySessionReplayRetentionsSecs to 0.

Optional: Override the default timeout values for all authentication sources.

Field	Description
Idle Timeout	Modify the default inactivity timeout value in the Idle Timeout field and select a unit of measurement from the list.
	You can enter an integer that represents a time period between 1 minute and 1,095 days. You can also empty the value to indicate that the inactivity timeout value should match the maximum lifetime.
	The default inactivity timeout value is 60 minutes.

- 3. Optional: Configure policy and settings for individual authentication sources under **Overrides**.
 - a. From the Authentication Source list, select an identity provider (IdP) adapter instance or an IdP connection.
 - b. Configure individual policy for the selected authentication source as follows.

Global policy (under Authentication Sessions)	Individual policy (under Overrides)	
The Enable Sessions for All Authentication Sources check box is not selected.	Select the Enable Sessions check box to enable authentication-session tracking for the selected authentication source.	
Authentication-session tracking is not enabled for all authentication sources.		
The Enable Sessions for All Authentication Sources check box is selected.	Clear the Enable Sessions check box to disable authentication-session tracking for the	
Authentication-session tracking is enabled for	selected authentication source.	
all authentication sources.	Select the Enable Sessions check box for the purpose of overriding other authentication-session settings for the selected authentication source.	



Note:

The **Enable Sessions** check box is not selected by default.

Keep in mind that for any HTML Form Adapter instance that has been configured to allow users to indicate whether their device is shared or private, if a user signs on without selecting the

This is my device check box on the login form, PingFederate removes authentication session information, if found, and does not store authentication sessions for the user.

c. Select the Persistent check box if your use cases require a longer session duration or a greater resilience against restarts of PingFederate and browsers.

Available and applicable only if the Enable Sessions check box is selected. The Persistent check box is not selected by default.



Note:

Persistent authentication sessions require an external storage.



Notes under step 2b apply here as well.

d. If authentication-session tracking is enabled for the selected authentication source and if you want to configure specific timeout values, select the Override Timeouts check box and configure timeout settings.

Field	Description
Idle Timeout	You can enter an integer that represents a time period between 1 minute and 1,095 days. You can also empty the value to indicate that the inactivity timeout value should match the maximum lifetime.
	This field has no default value.
Max Timeout	You can enter an integer that represents a time period between 1 minute and 1,095 days. You can also empty the value to indicate that the authentication sessions do not expire until they are removed from the system.
	The value of the Max Timeout field cannot be less than that of the Idle Timeout field.
	This field has no default value.
Unit	Select from the list the unit of measurement for both the Idle Timeout and Max Timeout fields.
	The default selection is Minutes .

- e. If authentication-session tracking is enabled for the selected authentication source and if you want to enforce authentication requirement based on the authentication context for the selected authentication source, select the Authentication Context Sensitive check box. This check box is not selected by default.
- f. Click Add.
- g. Repeat these steps to configure individual policy and settings for additional authentication sources.

Click Edit, Update, or Cancel to make or undo a change to an existing entry. Click Delete or **Undelete** to remove an existing entry or cancel the removal request.

4. To keep your configuration, click Save.

Result

When PingFederate authentication sessions are enabled, you can configure session-validation options for your OAuth use cases. These optional settings enable you to conjoin the validity of access tokens and the authentication sessions of the users. For more information, see Managing session validation settings on page 557.

Bundled adapters

The PingFederate Server comes bundled with a set of adapters for integration and customizable configuration, depending on your service needs.

PingFederate's bundled set of adapters:

- Composite adapter
- HTML Form adapter
- HTTP Basic adapter
- Identifier First adapter
- Kerberos adapter
- OpenToken adapter
- PingID adapter

Composite Adapter

PingFederate's Composite Adapter allows adapter chaining for single sign-on (SSO) requests that best fit your user or orgnization's needs.

For an Identity Provider (IdP), PingFederate includes a Composite Adapter that allows an administrator to chain the selection of available adapter instances for a connection. At runtime, adapter chaining means that SSO requests are passed sequentially through each adapter instance specified until one or more authentication results are found for the user.

Adapter chaining can be used to choose an adapter instance based on the method by which a user authenticated, or to integrate an organization's multi-factor authentication requirement.



For complex authentication requirements, consider implementing authentication policies. Go to Authentication # Policies. From the Policies window, make changes on the Policies tab.

Configuring a Composite Adapter instance

Configuring a Composite Adapter instance allows you to enable adapter chaining for a user authentication connection.

About this task

Configure an instance of a Composite Adapter in PingFederate using the administrative console.

Steps

- 1. Go to Authentication # Integration # IdP Adapters.
- On the IdP Adapters window, click Create New Instance to start the Create Adapter Instance configuration.

- 3. On the **Type** tab, configure the basics of this adapter instance.
 - a. Enter the instance name and ID.
 - b. From the **Type** list, select the adapter type.
 - c. Optional: From the **Parent Instance** list, select an existing type.
 If you are creating an instance that is similar to an existing instance, you might consider making it a child instance by specifying a parent. A child instance inherits the configuration of its parent unless overridden. You can specify overrides during the rest of the setup.
- 4. On the IdP Adapter tab, configure your Composite Adapter instance as follows:
 - a. Click Add a new row to 'Adapters'.
 - b. From the Adapter Instance list, select an IdP adapter instance. Click Update.

For more information, see the **Description** column in each configuration section and the following table.

PingFederate's column names and descriptions for creating an adapter instance

Column	Description
Policy (Required)	Required (the default) indicates authentication through this adapter instance is needed to continue SSO processing and to invoke any remaining instances in the chain. If you are integrating multifactor authentication, use this policy for each instance. The Composite Adapter instance returns an error when the authenticate attempt against a required adapter instance fails.
	Sufficient indicates that authentication through this adapter instance is enough to satisfy requirements, along with any required instances that have already been selected. Any subsequent configured instances in the chain are not invoked.
	Important: For the sufficient policy to work correctly, the adapter must return control to PingFederate after any kind of a failure.
AuthN Context Weight (Required)	If more than one adapter instance in the chain is capable of returning an authentication context, this relative weight is used to determine which value is included in the assertion, unless the value is overridden under AuthN Context Override .
	If weights are the same for two or more contexts, the first one processed is included in the assertion. The default value is 3.

c. Add at least one more adapter instance and configure its **Policy**, **AuthN Context Weight**, and **AuthN Context Override** settings.

Repeat this step to add more adapter instances as needed.

At runtime adapter chaining is sequential, starting at the top of the list.



Important:

You can configure several types of adapters -- for example, the IWA IdP Adapter -- to direct end users to an error page if authentication fails for any reason, which will halt further progress through a composite-adapter chain. For such adapter instances, ensure the **Error URL** option is not used in the instance configuration when continuation through an adapter-chaining sequence is required.

- d. Optional: In the **Action** column, manage the selected adapter instances.
- e. Configure Input User ID Mapping.

If you have configured any IdP Adapter developed using the IdpAuthenticationAdapterV2 interface from the PingFederate SDK, including the HTML Form Adapter, the Input User ID Mapping section appears. Additionally, some IdP adapters, such as the PingID Adapter and the separately available Symantec VIP Adapter, require a user ID to be passed in from an earlier-authentication step to perform multifactor authentication. If so, an administrator must specify the attribute containing the unique ID on this window. For example, to pre-populate the username

of an HTML Form Adapter instance with an attribute from an earlier authentication source in the previous steps:

- 1. Click Add a new row to 'Input User ID Mapping'.
- 2. From the **Target Adapter** list, select the HTML Form Adapter instance.
- 3. From the User ID Selection list, select a source attribute.
- 4. Click Update.



Note:

For OAuth use cases, entries in the Input User ID Mapping section might override the login hint parameter value provided by the OAuth clients when they submit their requests to the /as/authorization.oauth2 authorization endpoint.



By default, the HTML Form Adapter does not allow the users to change the username if it is configured to be pre-populated with an attribute from another authentication source. You can override this restriction by enabling the **Allow Username Edit** option on a per-adapter instance.

f. Configure Attribute Name Synonyms.

If any attributes are logically equivalent across two adapter instances but have different names. click Add a new row to 'Attribute Name Synonyms'. Select attributes from the Name and **Synonym** lists to create a mapping between them.

The attribute name under **Synonym** and its value are used in the SAML assertion when the two values returned from each adapter are identical. If returned values are different, both values are sent for the synonym.



Note:

Without this configuration to identify synonymous attribute names, both names and their values are sent in the SAML assertion.

g. Define the order in which different values are returned for the same attribute name.

For attributes of the same name configured in different adapter instances, you can change the order of returned values when the values are different. (Values are merged if they are the same.)

By default, the Add to Back value for an attribute name configured in the first instance is returned first and also listed first in the resulting SAML assertion. Then any different value from the same attribute name in a subsequently invoked instance is appended.

The order might not matter for many attributes, but in the case of the SAML-subject attribute, only the first value in the SAML assertion can be used for an SP connection partner under normal circumstances. Click **Add to Front** to reverse the default order, if needed.

5. On the Extended Contract tab, click Add to add attributes to be returned from each adapter instance configured on the previous window.



Note:

Attributes must correspond exactly to any or all of the attribute names listed on the Adapter Attribute tab for each configured adapter instance.



The Override Attributes check box in this window reflects the status of the override option in the Extended Contract tab.

a. Select the check box under Pseudonym for the user identifier of the adapter and optionally for the other attributes, if available.

This selection is used if any of your service provider (SP) partners use pseudonyms for account linking.



Note:

A selection is required whether or not you use pseudonyms for account linking. This allows account linking to be used later without having to delete and reconfigure the adapter. Ensure that you choose at least one attribute that is unique for each user, such as a user's email, to prevent assigning the same pseudonym to multiple users.

- b. Select the check box under Mask Log Values for any attributes that you want PingFederate to mask their values in its logs at runtime.
- c. Select the Mask all OGNL-expression generated log values check box, if OGNL expressions might be used to map derived values into outgoing assertions and you want those values masked.
- 7. Optional: On the Adapter Contract Mapping tab, configure the adapter contract for this instance with the following optional workflows:
 - Configure one or more data sources for datastore queries.
 - Fulfill adapter contract with values from the adapter, the default, datastore queries, if configured, context of the request, text, or expressions, if enabled.
 - Set up the Token Authorization framework to validate one or more criteria prior to the issuance of the adapter contract.
- 8. On the Summary tab, review your configuration and modify as needed. Click Done to exit the Create Adapter Instance configuration window.
- 9. When finished in the IdP Adapters window, click Save to confirm the adapter instance configuration. If you want to exit without saving the configuration, click Cancel.

HTML Form Adapter

The HTML Form Adapter supports user authentication when it occurs outside of the PingFederate server through an application or Identity Management system authentication module leveraging multiple user repositories and a password credential validator (PCV) instance.

Initial user authentication is normally handled outside of the PingFederate server using an application or an identity management system authentication module. Adapters or agents from PingFederate integration kits are typically used to integrate with these local authentication mechanisms.

PingFederate packages an HTML Form Adapter that delegates user authentication to a Password Credential Validator, such asan LDAP Username PCV. This authentication mechanism validates credentials against a user repository through an instance of a PCV. Multiple PCV instances can be added to an instance of the HTML Form Adapter to validate against multiple user repositories, in which case PingFederate falls to the subsequent PCV instance if the previous PCV instance fails to validate the user credentials.

When PingFederate receives an authentication request and the use case is associated with an HTML Form Adapter instance, PingFederate invokes the adapter if it does not find a valid authentication session. If customer IAM is configured and enabled, users can optionally register local accounts or sign on using third-party identity providers. If a user chooses to sign on using local accounts, the credentials are validated using the designated PCV instance or instances. If validated, PingFederate generates the requested SSO token or moves the request to the next checkpoint if authentication policies are involved.

In terms of the sign-on experience, the HTML Form Adapter allows you to use different customizable and localizable template files, define a logout path or a logout redirect page, notify users with password expiry information, allow users to change or reset their network passwords or redirect users to a company-hosted password management system, and enable self-service password reset, account unlock, and username recovery. All capabilities can be configured on a per-adapter instance basis.

PingFederate also tracks login attempts per adapter instance. This capability adds a layer of protection against brute force and dictionary attacks. When the Challenge Retries threshold is reached, the user is locked out for a period of time. The default value for the Challenge Retries setting is 3. If a higher value is preferred, consider reviewing the account lockout policy of the user repository first. For example, if the account lockout threshold is set to five on the target directory server and the Challenge Retries setting is also set to five or a higher value, the fifth sign-on attempt could potentially lock the user accounts on the directory server. The lockout period is controlled by the Account Locking Service on page 200.

This adapter does not provide an authentication context. For SAML connections, PingFederate sets the authentication context as follows:

- urn:oasis:names:tc:SAML:1.0:am:unspecified for SAML 1.x
- urn:oasis:names:tc:SAML:2.0:ac:classes:unspecified for SAML 2.0

As needed, the authentication context can be overridden by either an instance of the Requested AuthN Context Authentication Selector or the SAML AUTHN CTX attribute in the SAML attribute contract. The latter takes precedence.



Note:

The HTML Form Adapter is authentication API-capable. The PingFederate authentication API is a JSONbased API that enables end-user interactions, such as credential prompts, to be handled by an external web application. This API does so by providing access to the current state of the flow as an end user steps through a PingFederate authentication policy. For more information, see Authentication applications and the authentication API on page 391.

Configuring an HTML Form Adapter instance

Configure an HTML Form Adapter instance to validate a user authentication session with a Password Credential Validator (PCV) when your initial authentication needs to integrate with an external application or an identity management system (IdM) authentication module.

Steps

- 1. Go to Authentication # Integration # IdP Adapters.
- 2. On the IdP Adapters window, click Create New Instance to start the Create Adapter Instance configuration.
- 3. On the **Type** tab, configure the basics of this adapter instance.
 - a. Enter the instance name and ID.
 - b. From the **Type** list, select the adapter type.
 - c. Optional: From the Parent Instance list, select an existing type.

If you are creating an instance that is similar to an existing instance, you might consider making it a child instance by specifying a parent. A child instance inherits the configuration of its parent unless overridden. You can specify overrides during the rest of the setup.

- 4. On the IdP Adapter tab, configure your HTML Form Adapter instance as follows:
 - a. If you have not yet defined the desired Password Credential Validator instance, click Manage Password Credential Validators to do so.
 - b. Click Add a new row to 'Credential Validators' to select a credential-authentication mechanism instance for this adapter instance.
 - c. From the Password Credential Validator Instance list, select a Password Credential Validator instance. Click **Update**.

Add as many validators as necessary. Use the up and down arrows to adjust the order in which you want PingFederate to attempt credential authentication. If the first mechanism fails to validate the credentials, PingFederate moves sequentially through the list until credential validation succeeds. If none of the Password Credential Validator instances can authenticate the user's credentials, and the challenge retries maximum has been reached, the process fails.



Note:

If usernames overlap across multiple Password Credential Validator instances, this failover setup could lock out those accounts in their source locations.

d. Enter values for the adapter configuration, as described below.

Field	Description		
Challenge Retries	The account lockout threshold for this adapter instance. When the		
(Required)	number of login failures reaches this threshold, the user is locked out for a period time.		
	The default value is 3.		
Session State	Determines whether this HTML Form Adapter instance maintains adapter sessions and shares adapter sessions with other HTML Form Adapter instances.		
	Globally		
	Adapter sessions from this HTML Form Adapter instance are shared among other HTML Form Adapter instances that use the same Session State field value Globally .		
	Per Adapter		
	HTML Form Adapter maintains adapter sessions on a per-instance basis. Sessions from this HTML Form Adapter instance are not shared with other HTML Form Adapter instances.		
	None		
	This HTML Form Adapter does not maintain adapter sessions for this HTML Form Adapter instance.		
	Note:		
	To enable PingFederate authentication sessions globally or individually for this adapter instance, select None . For more information about PingFederate authentication sessions, see <i>Sessions</i> on page 276 and <i>Configuring authentication sessions</i> on page 279.		
	The default selection is None .		

Field

Description

Session Max Timeout

The maximum lifetime, in minutes, before an HTML Form Adapter session expires regardless of whether the Session Timeout field value has been reached. Ignored if **None** is selected for the **Session State** field.

Applicable only when the **Session State** field is set to **Globally** or **Per** Adapter.



Tip:

When you enable PingFederate authentication sessions globally or individually for this adapter instance, you can configure the Max **Timeout** setting for the same purpose. For more information, see Configuring authentication sessions on page 279.

The default value is 480 minutes, which translates to 8 hours.



This setting sets a maximum lifetime, subject to inactivity timeout. Consider the following examples:

- A user initiated an single sign-on (SSO) request at 9 a.m. and has not made another SSO request since then. At 10 a.m., the HTML Form Adapter session times out based on inactivity based on the default Session Timeout field value of 60 minutes.
- Another user initiated an SSO request at 9 a.m. and has been making SSO requests every hour at least once. This HTML Form Adapter session does not time out because the user has been actively making SSO requests; however, the HTML Form Adapter session does expire at 5 p.m. based on the default Session Max Timeout default value of 8 hours.
- If you leave both the Session Max Timeout and Session Timeout fields blank, HTML Form Adapter sessions do not expire until PingFederate restarts or the HTML Form Adapter sessions are cleaned up by another means.
- If you leave the **Session Max Timeout** field blank but set a value for the **Session Timeout** field, HTML Form Adapter sessions do not expire until they time out based on inactivity.



Session information is stored in the PF cookie. By default, the PF cookie is a session cookie and is typically removed when the user closes the browser.

You can optionally extend the lifetime of the PF cookie by editing the session-cookie-config.xml file, located in the <pf install>/ pingfederate/server/default/data/config-store directory. For more information, see Extending the lifetime of the PingFederate cookie on page 832.

Alternatively, you can enable PingFederate authentication sessions, store the authentication sessions externally, and leverage them as users request protected resources after restarting their browsers. For more information, see Sessions on page 276.

Field Description Allow Password Enables or disables the ability for users to change their network Changes password using this adapter instance as they initiate SSO requests and are prompted to enter their username and password. As needed, you can also provide your users the Change Password endpoint shown on the Summary window. The Change Password endpoint allows users to change their password without submitting SSO requests. For more information, see the /ext/pwdchange/Identify section in /ext/pwdchange/Identify on page 1137. Note: The LDAP Username Password Credential Validator (PCV) and the PingOne for Enterprise Directory PCV are currently the only PCVs bundled with PingFederate that support the change password feature. Important: When connecting to an Active Directory (AD) server, you must secure the datastore connection using LDAPS. AD requires this level of security to allow password changes. This check box is not selected by default. Password The URL for redirecting users to a company-specific password Management System management system to change their password. This field has no default value. Enable 'Remember My Allows users to store their username as a cookie when authenticating Username! with this adapter. Once stored, the username in the login form is prepopulated for subsequent transactions. Select the check box to enable the cookie functionality. Note: This option is hidden when users authenticate through a Composite Adapter instance that chains this adapter behind another authentication

source with an Input User ID Mapping configuration and the Allow Username Edits check box is not selected.

This check box is not selected by default.

Field Enable 'This is My Device' Change Password

Description

Allows users to indicate whether their device is shared or private. In this mode, PingFederate authentication sessions, if enabled, are not stored unless the user indicates the device is private. For more information about PingFederate authentication session, see Sessions on page 276.

This check box is not selected by default.



Note:

Adapter session tracking, if enabled by setting the Session State field to Globally or Per Adapter, is not affected by this configuration and the user's selection.

Notification

When selected, a notification is generated for the user who has successfully changed the password through the HTML Form Adapter. The destination is the user's email address, specifically the mail attribute value returned by the LDAP Username PCV instance.



Note:

This option requires the selection of the Allow Password Changes check box and a notification publisher instance. If you have not yet configured the desired notification publisher instance, click Manage **Notification Publishers.**

In addition, the LDAP Username PCV is the only PCV bundled with PingFederate that supports this notification feature.

This check box is not selected by default.

Show Password Expiring Warning

When selected, the HTML Form Adapter displays a warning to an authenticated user if the password associated with the account is about to expire soon. The message provides the number of days until the expiry of the current password and the options to change the password immediately or to snooze the message. Both the threshold and the snooze interval are configurable in the Advanced fields section; the default values are 7 days and 24 hours, respectively.



Note:

This option requires the selection of the Allow Password Changes check box. (Both check boxes are not selected by default.) In addition, the LDAP Username PCV is currently the only PCV bundled with PingFederate that supports the password expiring warning feature.

This check box is not selected by default.

Field Description

Password Reset Type

Select one of the following methods for self-service password reset (SSPR).

Authentication Policy

Based on the policy contract selected from the **Password Reset Policy Contract** list, PingFederate finds the applicable authentication policy to handle self-service password reset requests. If the users are able to fulfill the authentication requirements as specified by the policy, PingFederate allows the users to reset their password.

Email One-Time Link

Users receive a notification with a URL to reset their password. If you have not yet configured the desired notification publisher instance, click **Manage Notification Publishers**.

Email One-Time Password

Users receive a notification with a one-time password (OTP) to reset their password.

If you have not yet configured the desired notification publisher instance, click **Manage Notification Publishers**.

PingID

Users are prompted to follow the PingID authentication flow to reset their password.

Ensure the **PingID Username Attribute** field in the selected LDAP Username PCV instance is configured; otherwise, users will not be able to reset their password.

You must also download the settings file from the PingOne for Enterprise admin portal and upload the file to the **PingID Properties** advanced field.



Important:

Do not use a method that is already part of a multi-factor authentication policy that includes a password challenge, as that would indirectly reduce that authentication policy to a single factor. For example, if users normally authenticate with a password challenge and then PingID, the self-service password reset method should not be PingID. Instead, choose the **Authentication Policy** option, select a policy contract from the **Password Reset Policy Contract** list, and configure an authentication policy for self-service password reset.

Text Message

Users receive a text message notification with an OTP to reset their password.

Ensure the **SMS Attribute** field in the selected LDAP Username PCV instance is configured; otherwise, users will not receive text message notification for password reset.

If you have not yet configured SMS provider settings in PingFederate, click **Manage SMS Provider Settings**.

None

Users cannot reset password through this HTML Form Adapter instance.

Field

Description

Contract

Password Reset Policy If you use an authentication policy to handle SSPR requests, you must select a policy contract here.

> This policy contract doesn't require any extended attributes because uses this policy only to find the applicable authentication policies for password resets.



Important:

You must use a policy contract dedicated only to password reset. You can't use this policy contract for SSO anywhere else. To define a policy contract solely for password reset, click Manage Policy Contracts.

An authentication policy that uses this contract allows users to reset their password. The policy should use strong authentication methods to securely identify the user. To ensure that the user authenticating in the password reset flow is associated with the target account, you must map the incoming user ID into its authentication sources.

Account Unlock

Enables or disables the ability for users to unlock their account using this adapter instance as they initiate SSO requests and are prompted to enter their username and password.

As needed, you can also provide your users the Account Recovery endpoint shown on the Summary tab. The Account Recovery endpoint allows users to unlock their account without submitting SSO requests. For more information, see the /ext/pwdreset/Identify section in IdP endpoints.



Note:

You must also select a **Password Reset Type** value other than None and the Allow Password Changes check box as well because the initiating user must prove ownership of the account through the password reset flow.

Unlike self-service password reset self-service password reset (SSPR), when users succeed in proving account ownership, they are allowed to retain their current password or to reset their password as needed. Furthermore, self-service account unlock is only compatible with PingDirectory and Microsoft AD. If the underlying datastore is connected to an Oracle Unified Directory or Oracle Directory Server, users can only unlock their account by changing their current password through the password reset flow.

In addition, the LDAP Username PCV is the only PCV bundled with PingFederate that supports self-service account unlock.

This check box is not selected by default.

Local Identity Profile

Select a local identity profile to offer users the options to authenticate through third-party identity providers, self-register as part of the sign-on experience, and manage their accounts through a self-service profile management page.

There is no default selection.

- e. Optional: Click Show Advanced Fields to review or modify default values.
- f. If you have chosen **Text Message** as the password reset type, click **Manage SMS Provider Settings** at the bottom of the page to configure the SMS provider through which PingFederate can send text message notifications to the users.

5. On the **Extended Contract** tab, configure additional attributes for this adapter instance as needed.

The HTML Form Adapter contract includes two core attributes: username and policy.action. At runtime, PingFederate fulfills the policy.action core attribute as described in the following table.

Local identity profile	Runtime fulfillment
A selection is made.	If the local identity profile is configured with one or more authentication sources, and if the user chooses to register or authenticate with one of them, PingFederate sets the value to that authentication source. This design allows you to create rules in your authentication policies and form different policy paths for each authentication source. For more information, see <i>Enabling third-party identity providers</i> on page 341).
	Whether or not the local identity profile is configured with any authentication sources, if the user chooses to register directly by clicking on the Register now link, PingFederate sets the value to identity.registration. This fulfillment allows you to create rules to differentiate authentication requirements from the registration flow. For more information, see <i>Creating advanced registration mapping</i> on page 349.
No selection is made.	The policy.action attribute is not fulfilled.

6. On the Adapter Attributes tab, configure the pseudonym and masking options.



Note:

The **Override Attributes** check box in this window reflects the status of the override option in the Extended Contract tab.

a. Select the check box under Pseudonym for the user identifier of the adapter and optionally for the other attributes, if available.

This selection is used if any of your service provider (SP) partners use pseudonyms for account linking.



Note:

A selection is required whether or not you use pseudonyms for account linking. This allows account linking to be used later without having to delete and reconfigure the adapter. Ensure that you choose at least one attribute that is unique for each user, such as a user's email, to prevent assigning the same pseudonym to multiple users.

- b. Select the check box under Mask Log Values for any attributes that you want PingFederate to mask their values in its logs at runtime.
- c. Select the Mask all OGNL-expression generated log values check box, if OGNL expressions might be used to map derived values into outgoing assertions and you want those values masked.
- 7. Optional: On the Adapter Contract Mapping tab, configure the adapter contract for this instance with the following optional workflows:
 - Configure one or more data sources for datastore queries.
 - Fulfill adapter contract with values from the adapter, the default, datastore queries, if configured, context of the request, text, or expressions, if enabled.
 - Set up the Token Authorization framework to validate one or more criteria prior to the issuance of the adapter contract.
- 8. On the **Summary** tab, review your configuration and modify as needed. Click **Done** to exit the **Create** Adapter Instance configuration window.

HTML Form Adapter advanced fields

The following table describes the advanced fields that you can use to configure an HTML Form Adapter.

Advanced fields for setting password credentials and changes

Property	Description		
Login Template (Required)	The HTML template to prompt the users for their credentials. PingFederate allows each configured adapter instance to use a different login page template.		
(· · - ¶•··· • •)	The default template file is html.form.login.template.html.		
	Unless otherwise stated, all template files are located in the <pre>cpf_install>/</pre> <pre>pingfederate/server/default/conf/template directory.</pre>		
Logout Path	Any path in the format indicated. Setting a path invokes adapter logout functionality that is normally invoked during SAML 2.0 single-logout (SLO) processing. The resulting logout path is <code>/ext/<logout path=""></logout></code> . The logout path extends from the base URL. If virtual host names are configured, the logout path is accessible at those locations as well.		
	Available primarily for use cases where the partner software as a service (SaaS) providers who do not support SAML SLO but want the users' IdP SSO sessions to end after logging out of the SaaS services. For these use cases, the SaaS providers could redirect the users to the logout URL after the users log out of their platforms.		
	Note:		
	If specified, the path must be unique across all HTML Form Adapter instances, including child instances.		
	This field has no default value.		
Logout Redirect	The landing page at the service provider (SP) after successful identity provider (IdP) logout, applicable only when the Logout Path field is configured.		
	This field has no default value.		
Logout Template	The HTML template displayed when a user has successfully logged out in a configuration where the Logout Path field is configured but the Logout Redirect field is not.		
	The default template file is idp.logout.success.page.template.html.		
Change Password Template	The HTML template to prompt the users to change their password. PingFederate allows each configured adapter instance to use a different change password template.		
	The default template file is html.form.change.password.template.html.		
Change Password Message Template	The HTML template to be displayed when a user has successfully changed the password through the HTML Form Adapter.		
	The default template file is html.form.message.template.html.		

Property	Description		
Password Management System Message Template	The HTML template notifies the users that they are being redirected to a password management system to change their password.		
	The default template file is html.form.message.template.html.		
Change Password Email Template	The HTML email template PingFederate uses to generate the email message to notify the user that the password has been changed or reset successfully through the HTML Form Adapter.		
	The default template file is message-template-end-user-password-change.html, located in the $<\!pf_install\!>$ /pingfederate/server/default/conf/template/mail-notifications directory.		
	Applicable only if an instance of the SMTP Notification Publisher is selected from the Notification Publisher list.		
Expiring Password Warning Template	The HTML template to warn the users about approaching the password expiry day.		
	The default template file is <pre>html.form.password.expiring.notification.template.html.</pre>		
Threshold for Expiring Password Warning	The threshold, in days, to start warning the user about approaching the password expiry day.		
	The default value is 7 days.		
Snooze Interval for Expiring Password	The amount of time, in hours, to delay the next warning after the user has chosen to change the password later.		
Warning	The default value is 24 hours.		
Login Challenge Template	The HTML template to be displayed as the second step during a strong authentication. It is used to prompt the user to answer a challenge question after the first-factor login. The RADIUS Username Password Credential Validator (PCV) is an example of where it could be used.		
	The default template file is html.form.login.challenge.template.html.		
'Remember My Username' Lifetime	The number of days the cookie remains valid. Enter the number of days you want the username remembered in a cookie.		
	The cookie lifetime is reset upon each successful login in which the Remember my username check box on the login form is selected.		
	Note:		
	The value is ignored when users authenticate through a Composite Adapter instance that chains this adapter behind another authentication source with an Input User ID Mapping configuration and the Allow Username Edits check box is not selected.		
l	You can enter an integer between 1 and 3650.		
	The default value is 30 days.		

Property	Description		
'This is My Device' Lifetime	The number of days that a user's selection of the This is my device check box on the login form is retained.		
	The lifetime is reset upon each successful login in which the This is my device check box on the login form is selected.		
	You can enter an integer between 1 and 3650.		
	The default value is 30 days.		
Allow Username Edits During Chaining	When users authenticate through a Composite Adapter instance that chains this adapter behind another authentication source with an Input User ID Mapping configuration or initiate an OAuth authorization request with a login_hint parameter, the username in the login form is pre-populated. Users are not allowed to edit their usernames.		
	Select this check box if you want to allow users to edit the pre-populated username in the login form.		
	Note: Users who authenticate through a Composite Adapter instance without an Input User ID Mapping configuration or this adapter directly always need to enter their usernames.		
	This check box is not selected by default.		
Track Authentication Time	When selected, the time of authentication for each user is tracked and can be utilized by applicable use cases. For example, if an OAuth client sends an authorization request with a max_age parameter, such request will prompt the user to reauthenticate when the elapsed time between the current time and the time of the previous authentication is greater than the max_age value.		
	This check box is selected by default.		
Post-Password Change Re-Authentication Delay	The HTML Form Adapter reauthenticates the user using the new password immediately after a successful password change request. As needed, enter the amount of time, in milliseconds, that the adapter can wait prior to the reauthentication attempt.		
	The default value is 0, which is the minimum value. The maximum value is 60000, or 1 minute.		

Advanced fields for self-service password reset and account unlock

Property	Description	
Password Reset Username Template	The HTML template to prompt the user to enter a username for password reset.	
	This template applies for all password reset types other than None .	
	The default template file is forgot-password.html.	

Property	Description		
Password Reset Code Template	The HTML template to prompt the user to enter the one-time password (OTP) or the code for password reset.		
	This template applies when the password reset type is Email One-Time Password or Text Message .		
	The default template file is forgot-password-resume.html.		
Password Reset	The HTML template to prompt the user to define a new password.		
Template	This template applies for all password reset types other than None .		
	The default template file is forgot-password-change.html.		
Password Reset Error Template	The HTML template to notify the user that the password reset attempt has failed.		
	This template applies for all password reset types other than None .		
	The default template file is forgot-password-error.html.		
Password Reset Success Template	The HTML template to notify the user that the password reset attempt has succeeded.		
	This template applies for all password reset types other than None .		
	The default template file is forgot-password-success.html.		
Account Unlock Template	The HTML template to notify the user that the account unlock attempt has succeeded and to prompt the user to retain the current password or reset it.		
	The default template file is account-unlock.html.		
OTP Length	The number of characters in the one-time password for password reset.		
	The default value is 8.		
Password Reset Token	The validity in minutes for the one-time password or the one-time link.		
Validity Time	The default value is 10 minutes.		
PingID Properties	To configure self-service password reset using PingID, you must obtain the pingid.properties file and upload it to the HTML Form Adapter instance.		
	 Sign on to the PingOne admin portal. Go to Setup # PingID # Client Integration. Download the settings file pingid.properties. Close the PingOne admin portal. On the Manage IdP Adapters tab in the PingFederate administrative console, click Choose File. Select the pingid.properties file and click Open. 		

Property	Description	
Require Verified Email	When selected, requires that the user's email address is verified before sending a password reset, account unlock, or username recovery email.	
	If users are permitted to manage their accounts, they will be blocked from accessing any connected application until they have verified their email.	
	For more information about enabling user account management, see <i>Defining a local identity profile</i> on page 327.	
	By default, the check box is not selected.	
Username Recovery Template	The HTML template to prompt the user to enter an email address to recover the username associated with the account.	
	This template applies when username recovery is enabled.	
	The default template file is username.recovery.template.html.	
Username Recovery Info Template	The HTML template to notify the user to retrieve the email message with the recovered username.	
	This template is applicable when username recovery is enabled.	
	The default template file is username.recovery.info.template.html.	
Username Recovery Email Template	The HTML email template PingFederate uses to generate the email message containing the recovered username.	
	The default template file is message-template-username-recovery.html, located in the <pre>cpf_install</pre> /pingfederate/server/default/conf/template/mail-notifications directory.	
	Applicable only if an instance of the SMTP Notification Publisher is selected from the Notification Publisher list.	

CAPTCHA options

Property	Description	
CAPTCHA for Authentication	Enable CAPTCHA to protect the authentication process from automated attacks.	
CAPTCHA for Password Change	Enable CAPTCHA to protect the password change process from automated attacks.	
CAPTCHA for Password Reset	Enable CAPTCHA to protect the account recovery process for password reset and account unlock from automated attacks.	
CAPTCHA for Username Recovery	ne Enable CAPTCHA to protect the username recovery process from automat attacks.	

By default, CAPTCHA check boxes are not selected.

HTTP Basic Adapter

The HTTP Basic adapter provides user authentication through a password credential validator (PCV) to integrate PingFederate with local authentication mechanisms.

Initial user authentication is normally handled outside of the PingFederate server using an application or an Identity Management system authentication module. Adapters or agents from PingFederate integration kits are typically used to integrate with these local authentication mechanisms.

PingFederate packages an HTTP Basic Adapter that delegates user authentication to a *Password Credential Validator*, such as an LDAP Username PCV. This authentication mechanism validates credentials against a user repository through an instance of a PCV. Multiple PCV instances can be added to an instance of the HTTP Basic Adapter to validate against multiple user repositories, in which case PingFederate falls to the subsequent PCV instance if the previous PCV instance fails to validate the user credentials.

When PingFederate receives an authentication request and the use case is associated with an HTTP Basic Adapter instance, PingFederate invokes the adapter if it does not find a valid *authentication session*. If the HTTP Basic Adapter does not find a valid adapter session, it prompts the user for credentials.

This adapter does not provide an authentication context. For SAML connections, PingFederate sets the authentication context as follows:

- urn:oasis:names:tc:SAML:1.0:am:unspecified for SAML 1.x
- urn:oasis:names:tc:SAML:2.0:ac:classes:unspecified for SAML 2.0

As needed, the authentication context can be overridden by either an instance of the Requested AuthN Context Authentication Selector or the **SAML_AUTHN_CTX** attribute in the SAML attribute contract. The latter takes precedence.

Configuring an HTTP Basic Adapter instance

Configure an HTTP Basic Adapter instance to use credentials against a user repository through an instance of a password credential validator (PCV) to support user authentication when it occurs outside of the PingFederate server.

About this task

Using the administrative console, configure an HTTP Basic Adapter instance.

Steps

- 1. Go to Authentication # Integration # IdP Adapters.
- 2. On the IdP Adapters window, click Create New Instance to start the Create Adapter Instance configuration.
- 3. On the **Type** tab, configure the basics of this adapter instance.
 - a. Enter the instance name and ID.
 - b. From the **Type** list, select the adapter type.
 - c. Optional: From the **Parent Instance** list, select an existing type.

If you are creating an instance that is similar to an existing instance, you might consider making it a child instance by specifying a parent. A child instance inherits the configuration of its parent unless overridden. You can specify overrides during the rest of the setup.

- 4. On the **IdP Adapter** tab, configure your HTTP Basic Adapter instance as follows:
 - a. If you have not yet defined the desired Password Credential Validator instance, click Manage Password Credential Validators to do so.
 - b. Click Add a new row to 'Credential Validators' to select a credential-authentication mechanism instance for this adapter instance.
 - c. From the Password Credential Validator Instance list, select a Password Credential Validator instance. Click **Update**.

Add as many validators as necessary. Use the up and down arrows to adjust the order in which you want PingFederate to attempt credential authentication. If the first mechanism fails to validate the credentials, PingFederate moves sequentially through the list until credential validation succeeds. If none of the Password Credential Validator instances can authenticate the user's credentials, and the challenge retries maximum has been reached, the process fails.



Note:

If usernames overlap across multiple Password Credential Validator instances, this failover setup could lock out those accounts in their source locations.

d. Enter values for the adapter configuration.

See the on-window field descriptions and the following table for more information.

PingFederate's fields and descriptions for creating an HTTP Basic Adapter instance

Property	Description	
Realm	The name of a protected area. The value of this field is sent as a part of	
(Required)	the HTTP Basic authentication request. It appears in a dialog box that prompts the user for a username and password.	
	Note:	
	Once a user authenticates against a realm, if additional HTTP Basic Adapter instances share the same realm, the user is not prompted to reauthenticate.	
Challenge Retries (Required)	The number of attempts allowed for password authentication. The default value is 3.	

5. On the Extended Contract window, configure additional attributes for this adapter instance as needed.

The HTTP Basic Adapter contract includes one core attribute: username.

6. On the Adapter Attributes tab, configure the pseudonym and masking options.



The Override Attributes check box in this window reflects the status of the override option in the Extended Contract tab.

a. Select the check box under Pseudonym for the user identifier of the adapter and optionally for the other attributes, if available.

This selection is used if any of your service provider (SP) partners use pseudonyms for account linking.



Note:

A selection is required whether or not you use pseudonyms for account linking. This allows account linking to be used later without having to delete and reconfigure the adapter. Ensure that you choose at least one attribute that is unique for each user, such as a user's email, to prevent assigning the same pseudonym to multiple users.

- b. Select the check box under Mask Log Values for any attributes that you want PingFederate to mask their values in its logs at runtime.
- c. Select the Mask all OGNL-expression generated log values check box, if OGNL expressions might be used to map derived values into outgoing assertions and you want those values masked.
- 7. Optional: On the Adapter Contract Mapping tab, configure the adapter contract for this instance with the following optional workflows:
 - Configure one or more data sources for datastore queries.
 - Fulfill adapter contract with values from the adapter, the default, datastore queries, if configured, context of the request, text, or expressions, if enabled.
 - Set up the Token Authorization framework to validate one or more criteria prior to the issuance of the adapter contract.
- 8. On the **Summary** tab, review your configuration and modify as needed. Click **Done** to exit the **Create** Adapter Instance configuration window.
- 9. When finished in the IdP Adapters window, click Save to confirm the adapter instance configuration. If you want to exit without saving the configuration, click Cancel.

Identifier First Adapter

The Identifier First Adapter works best for use cases when a variety of user types are authenticating with PingFederate. The adapter analyzes the type of user and the credentials with which they have enrolled before, including datastore queries and user attributes, to provide support for user authentication.

When a variety of user types authenticate at PingFederate, it is often better to ask the user for their identifier first, determine their user population, and prompt the user with the desired authentication requirements and experience. The Identifier First Adapter is designed to handle this use case.

When PingFederate receives an authentication request and the use case is associated with an Identifier First Adapter instance, PingFederate invokes the adapter if it does not find a valid authentication session. The adapter prompts the user to enter their identifier and captures the identifier in the subject attribute.



Note:

subject is one of the two core attributes in the adapter contract. domain is the other one.

If the identifier is an email address, the adapter extracts the email address suffix and exposes it downstream through the domain attribute. Additionally, the adapter can leverage datastore queries to fulfill the domain attribute, or other extended attributes, to support identifiers of other kinds.

Based on the identification result and the configured authentication policies, PingFederate routes the user to the desired policy path. As the user fulfills the authentication requirements, the adapter preserves the identifier on the client side in a persistent cookie.

When the user signs off and makes a subsequent sign-on request from the same browser, the adapter offers the user to either select the previously authenticated identifier found in the cookie or to enter a new one. If the user opts to enter a new identifier, the adapter adds that identifier to the cookie once the user completes the authentication requirements.

The adapter keeps adding the most-recently-authenticated identifier until the number of identifier reaches a configurable limit. When the threshold is reached, the adapter removes the least-recently-used identifier from the cookie.

Lastly, the Identifier First Adapter also allow users to continue without entering or selecting an identifier, in which case it treats the authentication attempt as a failure and returns control to PingFederate. PingFederate can then route the request based on the configured policy path.

This adapter does not provide an authentication context. For SAML connections, PingFederate sets the authentication context as follows:

- urn:oasis:names:tc:SAML:1.0:am:unspecified for SAML 1.x
- urn:oasis:names:tc:SAML:2.0:ac:classes:unspecified for SAML 2.0

As needed, the authentication context can be overridden by either an instance of the Requested AuthN Context Authentication Selector or the SAML AUTHN CTX attribute in the SAML attribute contract. The latter takes precedence.



PingFederate creates authentication sessions configured for an Identifier First Adapter instance only when the complete single sign-on (SSO) transaction has succeeded. This lets the adapter prompt the user for a different user identifier when a chained adapter authentication fails because, for example, there's a typo in the user identifier.



Note:

The Identifier First Adapter is authentication API-capable. The PingFederate authentication API is a JSONbased API that enables end-user interactions, such as credential prompts, to be handled by an external web application. This API does so by providing access to the current state of the flow as an end user steps through a PingFederate authentication policy. For more information, see Authentication applications and the authentication API on page 391.

Configuring an Identifier First Adapter instance

Configure an instance of the Identifier First Adapter in PingFederate following these instructions and for additional configuration information on fieldnames, descriptions, and optimal settings depending on your use case.

About this task

Using the PingFederate administrative console, configure an Identifier First Adapter instance.

Steps

1. Go to Authentication # Integration # IdP Adapters.

- 2. On the IdP Adapters window, click Create New Instance to start the Create Adapter Instance configuration.
- 3. On the **Type** tab, configure the basics of this adapter instance.
 - a. Enter the instance name and ID.
 - b. From the **Type** list, select the adapter type.
 - c. Optional: From the **Parent Instance** list, select an existing type.
 If you are creating an instance that is similar to an existing instance, you might consider making it a child instance by specifying a parent. A child instance inherits the configuration of its parent unless overridden. You can specify overrides during the rest of the setup.
- 4. On the IdP Adapter tab, configure your Identifier First Adapter instance.

For more information about each field, see the following table.

PingFederate's fields and descriptions for creating an Identifier First Adapter instance

Field	Description	
Identifier Cookie Lifetime	Determines the number of days that previously authenticated identifiers are preserved as a cookie on the client side. This value can range from 0 through 3650.	
	Set to 0 to disable the storage of any previously authenticated identifiers.	
	The default value is 30.	
Allow Cancelling Identifier Selection	Determines whether a user is allowed to continue without entering or selecting an identifier.	
	If allowed, when a user decides to continue without providing an identifier, the Identifier First Adapter treats the authentication attempt as a failure and returns control to PingFederate.	
	This check box is not selected by default.	
Click Show Advanced	Fields to review the following settings. Modify as needed.	
Maximum Identifiers Count	Determines the maximum number of previously authenticated identifiers can be preserved in the identifier cookie. This value can range from 0 through 10 .	
	Set to $\ensuremath{\text{0}}$ to disable the storage of any previously authenticated identifiers.	
	The default value is 5.	
Identifier Selection Template	The HTML template to prompt the user to enter or select an identifier. PingFederate allows each configured adapter instance to use a different template as needed.	
	The default template file is identifier.first.template.html.	
	Like other Velocity template files, it is located in the <pf_install>/ pingfederate/server/default/conf/template directory.</pf_install>	

5. On the **Extended Contract** tab, configure additional attributes for this adapter instance as needed.

The Identifier First Adapter contract includes two core attributes: subject and domain.

If the identifier is an email address, the adapter extracts the email address suffix and exposes it downstream through the **domain** attribute. As needed, the adapter can leverage datastore queries to fulfill the **domain** attribute. For more information, see *step 7*).



The Override Attributes check box in this window reflects the status of the override option in the Extended Contract tab.

a. Select the check box under **Pseudonym** for the user identifier of the adapter and optionally for the other attributes, if available.

This selection is used if any of your service provider (SP) partners use pseudonyms for account linking.



Note:

A selection is required whether or not you use pseudonyms for account linking. This allows account linking to be used later without having to delete and reconfigure the adapter. Ensure that you choose at least one attribute that is unique for each user, such as a user's email, to prevent assigning the same pseudonym to multiple users.

- b. Select the check box under Mask Log Values for any attributes that you want PingFederate to mask their values in its logs at runtime.
- c. Select the Mask all OGNL-expression generated log values check box, if OGNL expressions might be used to map derived values into outgoing assertions and you want those values masked.
- 7. Optional: On the Adapter Contract Mapping tab, configure the adapter contract for this instance with the following optional workflows:
 - Configure one or more data sources for datastore queries.
 - Fulfill adapter contract with values from the adapter, the default, datastore queries, if configured, context of the request, text, or expressions, if enabled.
 - Set up the Token Authorization framework to validate one or more criteria prior to the issuance of the adapter contract.
- 8. On the **Summary** tab, review your configuration and modify as needed. Click **Done** to exit the **Create** Adapter Instance configuration window.
- 9. When finished in the IdP Adapters window, click Save to confirm the adapter instance configuration. If you want to exit without saving the configuration, click Cancel.

Identifier First Adapter and authentication policies

The Identifier First Adapter works best in conjunction with authentication policies and setting expected attribute values to enforce authentication requirements.

The Identifier First Adapter is designed to identify user populations. It supports email addresses natively: it extracts the email address suffix and exposes it downstream through the domain attribute. Additionally, the adapter can leverage datastore queries to fulfill the domain attribute or other extended attributes to support identifiers of other kinds.

The Identifier First Adapter is most effective when used in conjunction with authentication policies. Thepolicy paths are created by having rules matching expected values of the domain attribute or other extended attribute. Each expected value forms its own policy path, to which a series of authentication sources can be appended to enforce the desired authentication requirements.

For more information and configuration steps, see the subsequent sample use case.

Configure an Identifier First Adapter instance to determine user populations based on user identifiers usernames and an authentication policy to route sign-on requests to authentication sources tailored for their respective user populations.

About this task

Using the administrative console, follow the instructions below for configuring an Identifier First Adapter instance, creating an authentication policy to prompt the user for their identifier first, determining their user population, and routing the request to the desired authentication recommendations. Consider the sample use case here.

You need to enforce different sets of authentication requirements for two sets of users, employees, and external consultants.

Employees are given username@example.com email addresses, such as asmith@example.com. User records are stored in a local directory server. Employees sign on through an HTML Form Adapter instance.

Consultants have either username@example.org or username@example.info email addresses. User records are stored in a local database. Consultants can sign on using their username or email address and password through a local web portal. This web portal is integrated with PingFederate using the OpenToken framework.

Your organization owns another local database that keeps track of username, domain information, and email address for both employees and consultants. The column names are dsUid, dsDomain, and dsMail, respectively. For simplicity, no users share the same dsUid value.

In this sample use case, you must ensure that the Identifier First Adapter instance can handle the scenario where users may enter their email address or just their username when setting up the Identifier First Adapter instance. Additionally, when accessing protected resources, your organization has agreed to send the user's email address in the security token.

You have already created the following components:

- An LDAP datastore connecting to the local directory server. The attribute name of the user identifier is uid.
- An instance of the LDAP Username Password Credential Validator (PCV) validating credentials against the local directory server with the LDAP datastore. The LDAP Username PCV instance is extended with an additional attribute mail. The search filter is configured to handle identifiers in the format of an email address or a username. See the following code example.

```
(|(uid=${username}) (mail=${username}))
```

- An HTML Form Adapter instance delegating credential-validation to the LDAP Username PCV instance. The HTML Form Adapter instance is also extended with an additional attribute mail, which takes the mail attribute value from the LDAP Username PCV instance. The ID of this HTML Form Adapter instance is htmlForm.
- An OpenToken IdP Adapter instance digesting tokens from the web portal as the source of user attributes. The adapter contract is extended with an additional attribute mail. The web portal is designed to always include the user's email address in the token through the mail attribute. The ID of this OpenToken IdP Adapter instance is opentTokenIdp.

This sample use case requires the following additional components:

- An expression-enabled PingFederate environment. See step 1.
- An authentication policy contract to carry the email address from your organization to your partners.
 See <u>step 2</u>.
- A Java Database Connectivity (JDBC) datastore connecting to the database that hosts username, email, and domain information. See step 3.
- An Identifier First Adapter instance with an attribute source lookup configuration and a contract fulfillment via expressions for the domain adapter attribute. See step 4.

 An authentication policy to route user requests to different authentication sources based on user populations. See step 5.

To fulfill the requirements:

Steps

- 1. Enable expressions in PingFederate.
 - For configuration steps, see *Enabling and disabling expressions* on page 215.
- 2. Go to Authentication # Policies # Policy Contracts.
- 3. On the **Policy Contracts** window, click **Create New Contract** and create an authentication policy contract without any additional attributes.
- 4. Go to System # Data & Credential Stores # Data Stores.
- 5. From the **Data Stores** window, click **Add New Data Store**.
 - a. On the **Data Store Type** tab, enter **Name**.
 - b. From the **Type** list, select **Database (JDBC)** to create a JDBC datastore connection to the database that hosts username and domain information.

- 6. Create an instance of the Identifier First Adapter instance.
 - a. Follow steps 1 through 6 in Configuring an Identifier First Adapter instance on page 306.

For the sample use case, suppose you name the adapter instance ID 1st.

- b. Go to Applications # Integration # Adapter-to-Adapter Mappings.
- c. On the Adapter-to-Mappings window, select a source instance and a target instance, and click Add Mapping.
- d. On the Attribute Sources & User Lookup tab, click Add Attribute Source.



For more information about configuring the following steps, see *Datastore guery configuration* on page 362.

- 1. On the Data Store tab, enter an ID in Attribute Source ID and a name in Attribute Source Description, such as domainInfo and Domain Info, respectively. In the Active Data **Store**list, select the JDBC datastore created previously. Click **Next**.
- 2. On the Database Table and Columns tab, select the applicable options from the Schema and Table lists.
- 3. Under Columns to return from SELECT, select the dsDomain column and click Add Attribute. Click Next.
- 4. On the **Database Filter** tab, in the **Where** field, specify a filter to search by identifier that can handle identifiers in the format of an email address or a username.

See the following example of a filter entry.

```
dsUid='${subject}' OR dsMail='${subject}'.
```

- 5. Click Next.
- 6. On the **Summary** tab, click **Done**.
- e. On the Adapter Contract Fulfillment tab, configure as follows.

Contract	Source	Value
domain	Expression	
		<pre>#this.get("domain").toString().matches("(? i).+") ? #this.get("domain") : #this.get("ds.domainInfo.dsDomain")</pre>
		Note: Line breaks are inserted for readability only.
subject	Adapter	Not applicable. No selection is required.

The expression checks the domain attribute value returned by the Identifier First Adapter. If the value contains one or more character, PingFederate uses that as the value for the domain attribute. Otherwise, it uses the dsDomain column value returned from the JDBC datastore. In other words, this expression handles identifiers in the format of an email address or a username.

This sample expression is intended to demonstrate the capability of the Identifier First Adapter. Depending on the actual use cases, expressions may vary. For more information about expressions, see Construct OGNL expressions on page 216.

f. On the Issuance Criteria tab. click Next.



Note:

Depending on the actual use cases, you can add issuance criteria.

- g. On the Adapter-to-Adapter Summary tab, review your configuration instance. Click Done to save your adapter instance configuration.
- 7. Create an authentication policy with rules to form policy paths based on results from domain attribute values returned by the Identifier First Adapter.
 - a. Go to Authentication # Policies # Policies.
 - b. From the Policies tab, click Add Policy.
 - c. On the **Policy** window, enter a **Name**, and optionally a **Description**, for the policy.
 - d. From the **Policy** list, select the Identifier First Adapter instance created in step 5.
 - e. Click Rules to open the Rules dialog.
 - f. Add three rules as follows.

Defining Authentication Policy Rules dialog fields and entries

Attribute Name	Condition	Value	Result
domain	equal to	example.com	Example COM
domain	equal to	example.org	Example ORG
domain	equal to	example.info	Example INFO



Add one rule for each expected **domain** attribute value.

g. Clear the **Default to Success** check box to disable the option to specify a policy path for the scenario where the domain attribute value from the Identifier First Adapter instance does not match any configured value on the Rules dialog.

If you want to enable an authentication policy path for unexpected domain attribute values, leave the **Default to Success** check box as selected.

For more information about rules, see Configuring rules in authentication policies on page 246.

h. Click **Done** to close the **Rules** dialog.

Result:

By adding three rules and disabling the default to success option, the Identifier First Adapter instance now contains four policy paths: Fail, Example COM, Example ORG, and Example INFO.

i. Configure each policy path.

Fail

Select **Done**, which terminates the request in an error condition.

Example COM

Select the HTML Form Adapter instance, which contains two paths: Fail and Success.

Configure each policy path.

Fail

Select **Done**, which terminates the request in an error condition.

Success

Select the policy contract created in *step 2*.

Click **Options** to open the **Incoming User ID** dialog.

- 1. From the Sourcelist, select Adapter (ID 1st).
- 2. From the Attribute list, select subject.
- 3. Click **Done** to close the **Incoming User ID** dialog.

For more information, see Specifying an incoming user ID on page 245.

Example ORG (and then Example INFO)

Select the OpenToken IdP Adapter instance, which contains two paths: Fail and Success. Configure each policy path by using the same steps documented for the **Example COM** policy path

j. Configure contract fulfillment for each authentication policy contract as follows.

Contract Fulfillment fieldnames and entries

Result from rules	Contract Attribute	Source	Value
Example COM	subject	Adapter (htmlForm)	mail
Example ORG	subject	Adapter (openTokenIdp)	mail
Example INFO	subject	Adapter (openTokenIdp)	mail

For more information, see Configuring contract mapping on page 252.

k. Click Done. Click Save.

Result

You have now successfully configured an Identifier First Adapter instance and an authentication policy to prompt the user for their identifier first, determine their user population, and route the request to the desired authentication policy path.

Kerberos Adapter

The integrated Kerberos Adapter provides a seamless single sign-on (SSO) experience for Windows clients by authenticating SSO requests using the Kerberos v5 protocol against Active Directory (AD) domains.

When the PingFederate Identity Provider (IdP) server receives an authentication request for Service Provider-initiated SSO or a user clicks a hyperlink for IdP-initiated SSO, PingFederate invokes the Kerberos Adapter and returns to the browser an HTTP 401 Unauthorized response. When PingFederate receives a Kerberos ticket from the browser, it validates the ticket against the domain defined in the Kerberos Adapter configuration. If validation succeeds, PingFederate retrieves the username, the domain, and the security identifiers (SIDs) from the ticket; generates a SAML assertion with the username and optionally the associated domain, SIDs, or both; and passes it to the SP.



The Kerberos Adapter supports authentications by Kerberos only. If your environment requires NTLM support, you must deploy the IWA Integration Kit. You can safely deploy the IWA Adapter and create instances of it alongside with the Kerberos Adapter.

Authentication mechanism assurance

The integrated Kerberos Adapter supports authentication mechanism assurance from Active Directory domain service.

With an Identity Provider (IdP), you can use the Token Authorization framework to verify the SIDs value before issuing a token. Alternatively, you can map the SIDs value to an attribute in the contract and let the Service Provider (SP) determine if the user meets the requirements to access the protected resource. For the purpose of protecting resources based on sign-on method, authentication mechanism assurance from Active Directory (AD) domain service adds an additional group membership to the user's security identifiers attribute SIDs when a user signs on using a certificate-based sign-on method, such as a smart-card signon For example, you can restrict access to sensitive resources to users who sign on by using their smart cards, which requires a physical reader that you place in a physically secured location.

The integrated Kerberos Adapter supports authentication mechanism assurance by including the SIDs attribute of the authenticated user in the adapter contract.

If your use case requires authentication mechanism assurance, you can add a criterion in the Token Authorization framework to verify that the SIDs attribute contains the security identifier (SID) value associated with the required login method. If the SIDs attribute does not contain the specified SID value, the request is denied.



The SIDs attribute contains multiple values. Use the multi-value contains condition or the multi-value contains (case insensitive) condition to verify whether the SIDs attribute contains a specific value. You can also configure more complex evaluations using OGNL expressions.

Alternatively, you can map the SIDs attribute into the contract and let the SP determine if the user meets the requirements to access the protected resource.

For more information about authentication mechanism assurance, see the Authentication Mechanism Assurance for AD DS in Windows Server 2008 R2 Step-by-Step Guide from Microsoft's documentation.

Configuring a Kerberos Adapter instance for SSO authentication

An overview of creating and configuring a Kerberos Adapter instance to integrate PingFederate with Windows clients. Create and configure an instance of the Kerberos Adapter for Windows clients to authenticate using single sign-on.

Steps

- 1. Go to Authentication # Integration # IdP Adapters.
- On the IdP Adapters window, click Create New Instance to start the Create Adapter Instance configuration.
- 3. On the **Type** tab, configure the basics of this adapter instance.
 - a. Enter the instance name and ID.
 - b. From the **Type** list, select the adapter type.
 - c. Optional: From the Parent Instance list, select an existing type.

If you are creating an instance that is similar to an existing instance, you might consider making it a child instance by specifying a parent. A child instance inherits the configuration of its parent unless overridden. You can specify overrides during the rest of the setup.

4. On the **IdP Adapter** window, configure your Kerberos Adapter instance. See the on-window field descriptions and the following table for more information.

Field	Description	
Domain/Realm Name	Select your Windows domain.	
(Required)	If the domain or realm you want does not appear, click Manage Active Directory Domains/Kerberos Realms to add it. For more information, see <i>Active Directory and Kerberos</i> on page 915.	
Error URL Redirect	Enter a URL for redirecting the user if there are errors. This URL has an errorMessage query parameter appended to it, which contains a brief description of the error that occurred. The error page can optionally display this message on the window to provide guidance on remedying the problem.	
	Note:	
	In the case of an error, if you define an Error URL Redirect and the adapter instance is included in an instance of the Composite Adapter, the user is redirected to the configured error URL rather than continuing on to the next adapter in the chain. Leave this field blank to have the adapter continue on to the next adapter.	
	When employing the errorMessage query parameter in a custom error page, adhere to Web-application security best practices to guard against common content injection vulnerabilities. If no URL is specified, the appropriate default error landing page appears.	
Click Show Advanced I	Fields to review the following settings. Modify as needed.	
Error Template	When selected, displays a template to provide standardized information to the end user when authentication fails. The Error URL Redirect value is ignored.	
	The template <code>kerberos.error.template.html</code> in the <code><pf_install>/pingfederate/server/default/conf/template</pf_install></code> directory uses the Velocity template engine and can be modified in a text editor to suit your particular branding and informational needs. For example, you can give the user the option to try again if authentication fails. For more information on Velocity templates, see <code>Customizable user-facing pages</code> on page 798 .	
Authentication Context Value	This can be any value agreed to with your SP partner to indicate the type of credentials used to authenticate. Standard URIs are defined in the SAML specifications. For more information on SAML specifications, see the OASIS documents oasis-sstc-saml-core-1.1.pdf and saml-authn-context-2.0-os.pdf.	
	If left blank, PingFederate sets the authentication context as follows:	
	• urn:oasis:names:tc:SAML:1.0:am:unspecified for SAML 1. x	
	urn:oasis:names:tc:SAML:2.0:ac:classes:unspecified for SAML 2.0	
	Either an instance of the Requested AuthN Context Authentication Selector or the SAML_AUTHN_CTX attribute can override the authentication context in the SAML attribute contract. The latter takes precedence.	

- 5. On the **Extended Contract** tab, configure additional attributes for this adapter instance as needed. The Kerberos Adapter contract includes three core attributes: Domain/Realm Name, SIDs, and Username.
- 6. On the Adapter Attributes tab, configure the pseudonym and masking options.



Note:

The Override Attributes check box in this window reflects the status of the override option in the Extended Contract tab.

a. Select the check box under Pseudonym for the user identifier of the adapter and optionally for the other attributes, if available,

This selection is used if any of your service provider (SP) partners use pseudonyms for account linking.



Note:

A selection is required whether or not you use pseudonyms for account linking. This allows account linking to be used later without having to delete and reconfigure the adapter. Ensure that you choose at least one attribute that is unique for each user, such as a user's email, to prevent assigning the same pseudonym to multiple users.

- b. Select the check box under Mask Log Values for any attributes that you want PingFederate to mask their values in its logs at runtime.
- c. Select the Mask all OGNL-expression generated log values check box, if OGNL expressions might be used to map derived values into outgoing assertions and you want those values masked.
- 7. Optional: On the Adapter Contract Mapping tab, configure the adapter contract for this instance with the following optional workflows:
 - Configure one or more data sources for datastore queries.
 - Fulfill adapter contract with values from the adapter, the default, datastore queries, if configured, context of the request, text, or expressions, if enabled.
 - Set up the Token Authorization framework to validate one or more criteria prior to the issuance of the adapter contract.
- 8. On the Summary tab, review your configuration and modify as needed. Click Done to exit the Create Adapter Instance configuration window.
- 9. When finished in the IdP Adapters window, click Save to confirm the adapter instance configuration. If you want to exit without saving the configuration, click Cancel.

Configuring end-user browsers

Configure browsers at your site in order to use the Kerberos Adapter to authenticate users.

The client-side configuration requires the base URL or an applicable virtual host name of your PingFederate environment. Base URL is defined on the System # Server # Protocol Settings # Federation Info tab. To see a list of defined virtual host names, if configured, go to System # Server # Virtual Host Names.



Important:

If the browser is not properly configured, the user might be prompted to authenticate manually with their network credentials otherwise authentication fails the single sign-on (SSO) to the service providers.

Configuring Microsoft Internet Explorer

Configure Internet Explorer to support user authentication when using a Kerberos Adapter instance.

About this task

To configure Internet Explorer, edit the settings in Internet Options.

Steps

Add the base URL to Local intranet.



Note:

Skip this step if the base URL <pf-idp.domain.name> is internal and not fully qualified. For example, if the base URL is pingfederate, you can skip this step. However, if cpf-idp.domain.name is www.example.com, then you must add the base URL to the Sites list, as described in the following sub steps.

- a. Close all Internet Explorer tabs and windows.
- b. Open Control Panel # Internet Options.
- c. Click the Security tab.
- Select Local intranet, and then click Sites.
- e. Click Advanced.
- f. Enter the base URL. For example, www.example.com, and then click Add.
- g. Click Close. Click OK to return to the Security tab.
- 2. Verify **Automatic logon only in the Intranet zone** is selected.
 - Under the Security tab, select Local intranet and then click Custom level.
 - b. Verify Automatic logon only in the Intranet zone is selected in the Settings pane.
 - c. Click **OK** to return to the **Security** tab.
- 3. Verify proxy settings.



Note:

Skip the following sub steps if a proxy is not used.

- Click the Connections tab.
- b. Click LAN settings.
- c. Verify the Use a proxy server for your LAN ... check box is selected. Click Advanced.
- d. Enter the base URL in the Exceptions field, Click OK.
- e. Click OK to return to the Connections tab.
- 4. Verify **Enable Integrated Windows Authentication** is selected.
 - Click the Advanced tab.
 - b. Verify Enable Integrated Windows Authentication is selected in the Settings pane.
- 5. Click **OK** to close Internet Options.

Configuring Mozilla Firefox

Configure Firefox for Kerberos authentication at your site.

About this task

Configure Kerberos authentication using a Firefox browser following the steps below.

Steps

- 1. Start Firefox.
- 2. Open a new tab, and then enter about:config in the address bar.
- 3. Double-click the network.negotiate-auth.trusted-uris preference name to modify its value to include the base URL of your PingFederate environment. For example, www.example.com.
- 4. Click **OK** and close the about:config tab.
- 5. Optional: Exit Firefox.

OpenToken Adapter

In order to transfer identity and other user information between the PingFederate server and an end application, the PingFederate architecture allows for custom adapters to be deployed with the server.

PingFederate ships with a deployed OpenToken Adapter, which uses a secure token format OpenToken to transfer user attributes between an application and the PingFederate server.

On the identity provider (IdP) side, the OpenToken Adapter allows the PingFederate server to receive a user's identity from the IdP application.

For SAML connections, the IdP application can provide an authentication context to the service provider (SP) by including the authnContext attribute with the desired value in the secure token. Standard URIs are defined in the SAML specifications. For more information on assertions and protocol for SAML, see oasis-sstc-saml-core-1.1.pdf and saml-authn-context-2.0-os.pdf in the OASIS documentation.

If the secure token does not contain the authnContext attribute, PingFederate sets the authentication context as follows:

- urn:oasis:names:tc:SAML:1.0:am:unspecified for SAML 1.x
- urn:oasis:names:tc:SAML:2.0:ac:classes:unspecified for SAML 2.0

As needed, the authentication context can be overridden by either an instance of the Requested AuthN Context Authentication Selector or the SAML AUTHN CTX attribute in the SAML attribute contract. The latter takes precedence.

On the SP side, the OpenToken Adapter can be used to transfer user-identity information to the target SP application.

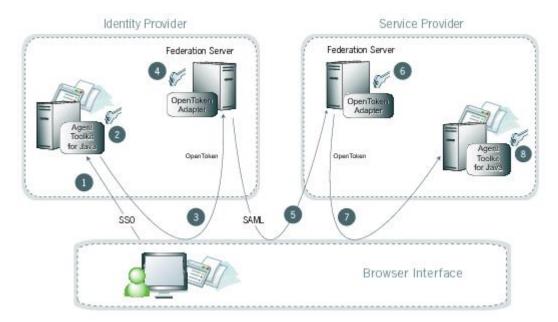
Specialized application integration kits are available from the Ping Identity Downloads website. Many kits leverage the OpenToken Adapter to integrate applications with the PingFederate server. The agent portions of the integration kits reside with the application and use the OpenToken to communicate with the OpenToken Adapter.



Note:

To integrate applications for use with the OpenToken Adapter, download an integration kit for PingFederate from the Ping Identity Downloads website and follow instructions for installing and using Agent Toolkits in the accompanying documentation. Follow the configuration instructions in Configuring an OpenToken IdP Adapter instance to setup and to use with your applications.

The following figure shows a basic IdP-initiated single sign-on (SSO) scenario using PingFederate with the Java Integration Kit on both sides of an identity federation.



IdP-Initiated SSO: POST/POST

Processing steps

- 1. A user initiates an SSO transaction.
- 2. The IdP application inserts attributes into the Agent Toolkit for Java, which encrypts the data internally and generates an OpenToken.
- 3. A request containing the OpenToken is redirected to the PingFederate IdP server.
- 4. The server invokes the OpenToken IdP Adapter, which retrieves the OpenToken, decrypts, parses, and passes it to the PingFederate IdP server. The PingFederate IdP server then generates a SAML assertion.
- 5. The SAML assertion is sent to the SP site.
- 6. The PingFederate SP server parses the SAML assertion and passes the user attributes to the OpenToken SP Adapter. The Adapter encrypts the data internally and generates an OpenToken.
- 7. A request containing the OpenToken is redirected to the SP application.
- 8. The Agent Toolkit for Java decrypts and parses the OpenToken and makes the attributes available to the SP Application.

Configuring an OpenToken IdP Adapter instance

Configure an instance of the OpenToken IdP Adapter in PingFederate.

About this task

Configure an OpenToken Identity Provider (IdP) Adapter instance using the administrative console to enable a secure authentication plugin for your custom application.

Steps

- 1. Go to Authentication # Integration # IdP Adapters.
- 2. On the IdP Adapters window, click Create New Instance to start the Create Adapter Instance configuration.

- 3. On the **Type** tab, configure the basics of this adapter instance.
 - a. Enter the instance name and ID.
 - b. From the **Type** list, select the adapter type.
 - c. Optional: From the Parent Instance list, select an existing type. If you are creating an instance that is similar to an existing instance, you might consider making it a child instance by specifying a parent. A child instance inherits the configuration of its parent unless overridden. You can specify overrides during the rest of the setup.
- 4. On the IdP Adapter tab, configure your OpenToken IdP Adapter instance.



These values depend on your developer's implementation.

For more information, see the **Description**field provided on-window and inthe following table.

PingFederate's field names and descriptions for creating an OpenToken IdP Adapter instance

Field	Description	
Password	The password to use for generating the encryption key. It is also known a	
Confirm Password	the shared secret.	
(Required)		
Authentication Service	The URL to which the user is redirected for a single sign-on (SSO)	
(Required)	event. This URL is part of an external application, which performs user authentication.	
Click Show Advanced Modify as needed.	Fields in the Instance Configuration tab to review the following settings.	
Transport Mode	How the token is transported to and from the application, either through a query parameter, a cookie (default), or as a form POST.	
Token Name	The name of the cookie or query parameter that contains the token. This	
(Required)	name must be unique for each adapter instance. Override the default value opentoken as needed.	
Cipher Suite	The algorithm, cipher mode, and key size that should be used for encrypting the token. The default selected value is AES-128/CBC .	
Logout Service	The URL to which the user is redirected for a single-logout event. This URL is part of an external application, which terminates the user session.	
Cookie Domain	The server domain; for example, example.com. If no domain is specified, the value is obtained from the request.	
Cookie Path	The path for the cookie that contains the token.	
Token Lifetime	The duration in seconds for which the token is valid. Valid range is 1 to	
(Required)	28800. The default value is 300 (5 minutes).	
Session Lifetime	The duration in seconds for which the token may be re-issued without	
(Required)	authentication. Valid range is 1 to 259200. The default value is 43200 (12 hours).	
Not Before Tolerance	The amount of time in seconds to allow for clock skew between servers. Valid range is 0 to 3600. The default value is \circ .	
(Required)		

Field	Description	
Force SunJCE Provider	If selected, the SunJCE provider is forced for encryption and decryption.	
Use Verbose Error Messages	If selected, use verbose TokenException messages.	
Obfuscate Password	If selected, the default, the password is obfuscated and password-strength validation is applied. Clearing the check box allows backward compatibility with previous OpenToken agents.	
Session Cookie	If selected, OpenToken is set as a session cookie, rather than a persistent cookie. Applies only if the Transport Mode field is set as Cookie . The check box is not selected by default.	
Secure Cookie	If selected, the OpenToken cookie is set only if the request is on a secure channel (https). Applies only if the Transport Mode field is set to Cookie . The check box is not selected by default.	
Delete Cookie	If selected, the token cookie is deleted immediately after consumption. Applies only if the Transport Mode field is set to Cookie . The check box is not selected by default.	
Replay Prevention	Selecting this option is recommended only if Query Parameter is the chosen token transport mode and form POST is used by an associated connection to send the SAML assertion. If selected, PingFederate ensures that the token can be used only once. By default, the check box is not selected.	
	Note:	
	Selecting this option might affect resource utilization and performance.	
Skip Malformed Attribute Detection	If not selected, the default, it prevents insecure content from affecting the security of your application and the agent. Update your applications with the latest version of the agent. We recommend not to change the value of this flag.	

5. On the **Actions** tab, click **Download** under **Action Invocation Link**, and then click **Export** to save the properties file.

The values in the resulting file, <code>agent-config.txt</code>, represent the console configuration and are used by the Identity Provider (IdP) application. See the documentation of your respective integration kit for more information.

6. On the Extended Contract tab, configure additional attributes for this adapter instance as needed.

The OpenToken IdP Adapter contract includes one core attribute: subject.

The OpenToken IdP Adapter always extends the core contract with an attribute userId as well and fulfills it with the value of subject for backward compatibility reason.



The Override Attributes check box in this window reflects the status of the override option in the Extended Contract tab.

a. Select the check box under **Pseudonym** for the user identifier of the adapter and optionally for the other attributes, if available.

This selection is used if any of your service provider (SP) partners use pseudonyms for account linking.



Note:

A selection is required whether or not you use pseudonyms for account linking. This allows account linking to be used later without having to delete and reconfigure the adapter. Ensure that you choose at least one attribute that is unique for each user, such as a user's email, to prevent assigning the same pseudonym to multiple users.

- b. Select the check box under Mask Log Values for any attributes that you want PingFederate to mask their values in its logs at runtime.
- c. Select the Mask all OGNL-expression generated log values check box, if OGNL expressions might be used to map derived values into outgoing assertions and you want those values masked.
- 8. Optional: On the Adapter Contract Mapping tab, configure the adapter contract for this instance with the following optional workflows:
 - Configure one or more data sources for datastore queries.
 - Fulfill adapter contract with values from the adapter, the default, datastore queries, if configured, context of the request, text, or expressions, if enabled.
 - Set up the Token Authorization framework to validate one or more criteria prior to the issuance of the adapter contract.
- 9. On the **Summary** tab, review your configuration and modify as needed. Click **Done** to exit the **Create** Adapter Instance configuration window.
- 10. When finished in the **IdP Adapters** window, click **Save** to confirm the adapter instance configuration. If you want to exit without saving the configuration, click Cancel.

Configuring an OpenToken SP Adapter instance

Configure an instance of the deployed OpenToken Adapter, which uses a secure token format to transfer user attributes between an application and the PingFederate server.

About this task

Configure an OpenToken Service Provider (SP) Adapter instance to enable a secure transfer of the useridentity information to the target SP application.

Steps

- 1. Go to Applications # Integration # SP Adapters to access the Manage SP Adapters Instances
- Click Create New Instance to start the Create Adapter Instance configuration wizard.

- 3. On the **Type** tab, configure the basics of this adapter instance.
 - a. Enter the Instance Name, Instance ID, and Parent Instance information and select the adapter type from the Type list.
 - b. Optional: Select a Parent Instance from the list.

This is useful when you are creating an instance that is similar to an existing instance. The child instance inherits the configuration of its parent. In addition, you have the option to override one or more settings during the rest of the setup. Select the Override ... check box and make the adjustments as needed in one or more subsequent windows.

4. On the Instance Configuration tab, configure your OpenToken SP Adapter instance security context.



These values are dependent on your developer's implementation.

For more information, see the **Description** field provided in-window and in the following table.

PingFederate's field names and descriptions for creating an adapter instance

Field	Description	
Password	The password to use for generating the encryption key. It is also known as the shared secret.	
Confirm Password		
(Required)		
Click Show Advanced Modify as needed.	Fields in the Instance Configuration tab to review the following settings.	
Transport Mode	How the token is transported to and from the application, either through a query parameter, a cookie, or as a form POST (default).	
Token Name	The name of the cookie or query parameter that contains the token. This name must be unique for each adapter instance. Override the default value opentoken as needed.	
Cipher Suite	The algorithm, cipher mode, and key size that should be used for encrypting the token. The default selected value is AES-128/CBC .	
Authentication Service	The URL to which the user is redirected for a single sign-on (SSO) event. This URL overrides the Target Resource, which is sent as a parameter to the Authentication Service.	
Account Link Service	The URL to which the user is redirected for account linking. This URL is part of an external SP application. This external application performs user authentication and returns the local user ID inside the token.	
Logout Service	The URL to which the user is redirected for a single-logout event. This URL is part of an external application, which terminates the user session.	
Cookie Domain	The server domain; for example, ${\tt example.com}.$ If no domain is specified, the value is obtained from the request.	
Cookie Path	The path for the cookie that contains the token.	
Token Lifetime (Required)	The duration in seconds for which the token is valid. Valid range is 1 to 28800. The default value is 300 (5 minutes).	

P'.1.1	December 1997	
Field	Description	
Session Lifetime	The duration in seconds for which the token may be re-issued without authentication. Valid range is 1 to 259200. The default value is 43200, 12 hours.	
(Required)		
Not Before Tolerance	The amount of time in seconds to allow for clock skew between servers. Valid range is 0 to 3600. The default value is 0.	
(Required)		
Force SunJCE Provider	If selected, the SunJCE provider is forced for encryption/decryption.	
Use Verbose Error Messages	If selected, use verbose TokenException messages.	
Obfuscate Password	If selected, the default, the password is obfuscated and password-strength validation is applied. Clearing the check box allows backward compatibility with previous OpenToken agents.	
Session Cookie	If selected, OpenToken is set as a session cookie rather than a persistent cookie. Applies only if the Transport Mode field is set to Cookie . The check box is not selected by default.	
Secure Cookie	If selected, the OpenToken cookie is set only if the request is on a secure channel (https). Applies only if the Transport Mode field is set to Cookie . The check box is not selected by default.	
Send Subject as Query Parameter	Selecting this check box sends the user identifier subject as a clear-text query parameter, if the Transport Mode field is set to Query Parameter . If Form POST is the chosen token transport mode, the user identifier is sent as POST data.	
Subject Query Parameter	The parameter name used for the user identifier when the Send Subject ID as Query Parameter check box is selected.	
Send Extended Attributes	Extended Attributes are typically sent only within the token, but this option overrides the normal behavior and allows the attributes to be included in browser cookies or query parameters.	
Skip Trimming of Trailing Backslashes	If not checked, the default, it prevents insecure content from affecting the security of your application and the agent. Update your applications with the latest version of the agent. We recommend not to change the value of this flag.	
URL Encode Cookie Values	If checked, the extended attribute cookie value will be URL encoded.	

5. In the **Actions** tab, click **Download** under **Action** section. Click **Export** to save the properties file.

The values in the resulting file, agent-config.txt, represent the console configuration and are used by the SP application. See the documentation of your respective integration kit for more information.

- 6. Optional: In the Extended Contract tab, configure additional attributes for this adapter instance.
- 7. In the **Summary** tab, review your configuration, modify as needed. Click **Done**.
- 8. On the SP Adapters window, click Save to confirm the adapter instance configuration.

If you want to exit without saving the configuration, click Cancel.

Customer IAM configuration

PingFederate empowers administrators to deliver a secure and easy-to-use customer authentication, registration, and profile management solution. This solution leverages the HTML Form Adapter to offer

users the options to authenticate through third-party identity providers (IdPs), self-register as part of the sign-on experience, and manage their accounts through a self-service profile management page.

Like other user-facing windows, administrators can customize and localize both the registration and profile management pages to present a consistent branding experience based on the needs of the users and the organizations.

Furthermore, administrators can allow users to leverage their existing identities from third-party identity providers. Any IdP connection or IdP adapter, such as the LinkedIn Cloud Identity Connector, can be used as an authentication source to a third-party identity provider. This optional capability enables a mapping configuration between the attributes returned by the identity provider and the fields within the registration page, streamlining the registration process.

Depending on the requirements and configuration of existing components, the configuration process might involve changes to these configuration components: authentication policy contracts, local identity profiles, HTML Form Adapter instances, and IdP authentication policies.



The HTML Form Adapter is authentication API-capable. The PingFederate authentication API is a JSONbased API that enables end-user interactions, such as credential prompts, to be handled by an external web application. This API does so by providing access to the current state of the flow as an end user steps through a PingFederate authentication policy. For more information, see Authentication applications and the authentication API on page 391.

Setting up PingDirectory for customer identities

PingFederate can optionally store customer identities in PingDirectory. After you have installed PingDirectory, update the LDAP schema with a new object class and a couple attributes to store customer identities and their connections.

About this task

Update the LDAP schema with a new object class and a couple attributes using an LDIF file provided. To optimize performance, apply updates in indexes to the directory as well. In addition, you must configure in PingFederate an LDAP datastore connection to your PingDirectory and an LDAP Username Password Credential Validator instance for the HTML Form Adapter to validate user credentials. If you have previously created these components, you can reuse them.



Note:

Skip this configuration if your use case does not involve registration or profile management. For more information, see Enabling third-party identity providers without registration on page 352.

Steps

- 1. Update the LDAP schema.
 - a. Sign on to the PingDirectory administrative console.
 - b. Go to the LDAP Schema # Schema Utilities screen.
 - c. Click Import Schema Element.
 - d. Copy the schema changes from the <pf install>/pingfederate/server/default/ paste them into the text area.
 - If you are creating a new organizational unit as part of the LDIF import, edit the DN information.
 - e. Click Import.

Use PingDirectory's dsconfig utility to create this index. The dsconfig utility is interactive. You can also provide inputs as command arguments. For example, the following samples create the pfconnected-identity index.

```
$ bin/dsconfig create-local-db-index \
                  --backend-name userRoot \
                  --index-name pf-connected-identity \
                  --set index-type:equality
```

After adding the index, use the rebuild-index utility to build the indexes. For instance, the following sample builds the required index.

```
$ bin/rebuild-index \
                           --baseDN "dc=example,dc=com" \
                           --index pf-connected-identity
```

3. Create an LDAP datastore connection to your PingDirectory on System # Data Stores.

If you have already created an LDAP datastore connection to your PingDirectory, you can reuse it.

4. Create an instance of the LDAP Username Password Credential Validator on System # Password **Credential Validators** to validate user credentials stored in PingDirectory.

If you have already created an LDAP Username Password Credential Validator instance, you can reuse it.



Later you will create a local identity profile as part of the customer IAM configuration. The Search Base value here should match the Base DN value defined in the local identity profile. For more information, see Configuring LDAP base DN and attributes on page 334.

Managing local identity profiles

When associated with an HTML Form Adapter instance, a local identity profile provides users the option to authenticate through third-party identity providers, self-register as part of the sign-on experience, and manage their accounts through a self-service profile management page.

About this task

A typical customer identity and access management (CIAM) use case only requires one local identity profile. As needed, you can create multiple profiles to suit the needs of your organization. Using the administrative console. Local identity profiles are defined in the **Identity Policies** section.



Note:

As of PingFederate 10.1, an authentication session is automatically created for a user after registration, preventing the user from having to log in again during the next single sign-on (SSO) transaction. This feature is enabled by default for all new and existing local identity profiles. However, if needed, you can disable it through the /localIdentity/identityProfiles administrative API endpoint by setting the createAuthnSessionAfterRegistration attribute to false.

Steps

 To configure a new profile, go to Authentication # Policies # Local Identity Profiles. Click Create **New Profile.**

- To modify an existing profile, select it by its name under Local Identity Profile Name.
- To review the usage of an existing profile, click Check Usage under Action.
- To remove an existing instance or to cancel the removal request, click **Delete** or **Undelete** under **Action**.

Defining a local identity profile

Defining a local identity profile provides users the option to authenticate with third-party identity providers, to self-register for sign-on, and to self-service manage their profile.

About this task

Using the administrative console, define a local identity profile and choose an authentication policy contract, which is the set of attributes that are returned from an authentication policy that uses this profile. In addition, choose whether to enable self-service registration and profile management.

Steps

- 1. Go to Authentication # Policies # Local Identity Profiles, and then click Create New Profile.
- 2. On the Profile Info tab, enter a name in the Local Identity Profile Name field.
- From the Authentication Policy Contract list, select a contract.
 If you have not yet defined the desired contract, click Manage Policy Contracts.
- 4. Select the Enable Registration check box if you want to enable users to complete a self-service registration as part of the sign-on experience through an instance of the HTML Form Adapter. This check box is not selected by default.
- Select the Enable Profile Management check box if you want to allow users to manage their accounts.
 - This check box is not selected by default.
- 6. When finished, click Next.

Defining authentication sources

Authentication sources are identifiers for third-party identity providers, such as social providers used to display these providers on the HTML form adapter user interface as alternate authentication and registration options. They are also used in authentication policies to configure branches to identity provider (IdP) adapters and connections.

About this task

Using the administrative console, configure a new authentication source, modify existing ones, and configure storage settings for attributes as needed.

Authentication sources are optional. They are the identifiers for third-party identity providers, such as social network providers. When defined, the associated HTML Form Adapter instance displays them on the signon page as alternative options for authentication and registration, if enabled. If profile management is enabled, users can connect or disconnect third-party identity providers to and from their accounts.

You can store attributes received from third-party identity providers as part of the user records. If required, attributes can be updated as users authenticate. By default, attributes are removed from user records as users disconnect third-party identity providers from their accounts. It is worth noting that storing attributes received from third-party identity providers is optional and configurable on a per-local identity profile basis. Additionally, this option is only applicable when a local identity profile is configured with registration, profile management, or both.

Steps

1. Configure authentication sources.

Action	Steps
Add a new authentication source	 a. Go to Authentication # Policies # Local Identity Profiles. b. On the Local Identity Profiles window, click Create New Profile. c. On the Profile Info tab, enter Local Identity Profile Name, from the Authentication Policy Contract list, select an option. d. Select the Enable Registration and Enable Profile Management check boxes as needed.
	Tip: If you use the authentication source names Facebook, Google, LinkedIn, Twitter, or FIDO, the HTML Form Adapter default templates render the associated icons on the registration and profile management pages.
Modify an existing authentication source	 a. On the Authentication Sources tab, under Action, click Edit. b. Click Update or Cancel.
Remove an existing authentication source	 a. On the Authentication Sources tab, under Action, click Update for the applicable authentication source. b. Click Delete.
	When removing an authentication source, keep in mind that accounts that were created using the associated third-party identity provider will no longer be usable after the removal. To minimize the risk of accidental removals, the administrative console prompts to confirm each removal request.
Change the display order of the authentication sources on the sign-on page and the profile management page	Use the up and down arrows to reorder them.

Result:

Make a note of the values defined here. In a later step, you will create a rule for each authentication source in an identity provider (IdP) authentication policy. Each rule forms a policy path that initiates the authentication process.

2. Configure storage settings for attributes received from third-party IdP.



The attribute storage settings are inapplicable and not shown if neither Enable Registration nor Enable Profile Management check box not selected on the Profile Info tab.

Choose from:

- If storing attributes, on the Authentication Sources tab, select the Store Attributes check box. This check box is not selected by default.
- If you want attributes retained after users disconnect third-party IdP from their accounts, select the Keep Attributes After Users Disconnect check box.

This check box is not selected by default.

 If you want attributes updated as users authenticate, select the Update Attributes When Users Authenticate check box and enter a value in the Minimum Number of Days Between Updates field.

The Update Attributes When Users Authenticate check box is not selected by default, and the Minimum Number of Days Between Updates field has no default value.

3. When finished, click Next.

Defining local identity fields

Configuring local identity fields determines which fields will be displayed as input controls on the user registration and profile management pages.

About this task

Using the administrative console, define the local identity fields that suit your registration and profile management requirements.

When registration is enabled for a local identity profile, select a local identity field to be the unique identifier for the purpose of identifying the users. To enable email ownership verification, add a field to store the email address and another field to store the verification status; while the former can be any field that uses the **Email** or **Text** input control, the latter must use the **Hidden** input control.

Steps

To access the Policies section in the administrative console, go to Authentication # Policies # Local Identity Profiles. This will open the Local Identity Profiles window configuration.

Action	Steps
Access the Policies section in the administrative console	Go to Authentication # Policies # Local Identity Profiles. This will open the Local Identity Profiles window configuration.
Add new local identity fields	In the Local Identity Profiles window, from the Local Identify Profile Name section, select the applicable identity profile. On the Fields tab, click Create New Field.
Select one of the local identity fields as the unique identifier	On the Fields tab, select the Unique ID option for the applicable field. Applicable and required only if registration is
	enabled on the Profile Info tab.

Action	Steps
	Any field that uses the Checkbox, Checkbox Group, Date, or Dropdown input control cannot be chosen as the unique identifier because values from such field will likely collide as the population of users grows.
Modify an existing local identity field	On the Fields tab, click Edit for the applicable field.
Remove an existing local identity field or to cancel the removal request	On the Fields tab, click Delete or Undelete for the applicable field.
Change the display order of the local identity fields on the registration page and the profile management page	On the Fields tab, use the up and down arrows to reorder them.
Mask local identity field values in logs for the configuration scenario where OGNL expressions might be used to map derived values into outbound single sign-on (SSO) tokens in authentication policies	On the Fields tab, select the Mask all OGNL-expression generated log values check box.

Configuring a local identity field

Configure a local identity field to determine which fields will be displayed as input controls on the user registration and profile management pages.

About this task

Using the administrative console, create a new or modify an existing local identity field.

Steps

- 1. Go to Authentication # Policies # Local Identity Profiles.
- 2. In the Fields tab, click Create New Field.

Result: This will open the Field Configuration window and tab.

3. From the **Type** list, select an option.

You cannot change the identifier of an existing field.

- 4. Enter an ID in the ID field.
- 5. Enter a name in the Label field.

This is the field name that users see on the registration and profile management pages.

- 6. Under **Type**, select one of the following input controls from the list.
 - Checkbox
 - Checkbox Group
 - Date
 - Dropdown
 - Email
 - Phone
 - Text
 - Hidden

7. To configure this field to appear on the registration or profile management page, under the **Applies To** section, select **Profile Management**, or **Registration** check boxes, or both.

This step is applicable only if both **Enable Registration** and **Enable Profile Management** are enabled on the **Profile Info** tab. Both pages are selected by default.

If only either registration or profile management is enabled., all fields, with the exception of hidden fields, are shown on the registration or the profile management tab.

8. Optional: Select the relevant parameters under **Parameters**.

Example:

You can make a non-hidden field mandatory or read-only. You can also configure PingFederate not to record values from this field in logs.

9. Optional: Enter a value under **Default Value**.

Specifying a default value can streamline the registration process. This is the default value of the field unless another value is specified in the authentication policy. For more information, see *Configuring local identity mapping* on page 252.

Default Value is not shown if you have chosen an input control of **Checkbox group**, **Email**, **Phone**, or **Hidden**, or the **Read-Only** parameter.

10. Add the applicable predefined values under **Options**.

This step is applicable and required only if you have chosen **Checkbox Group** or **Dropdown** as the input control.

11. Click Done.

Result:

The administrative console brings back the **Field** window, where you can configure other options and save your changes.

Configuring email ownership verification options

Based on your customer IAM use cases, you can optionally offer users the opportunity to confirm the ownership of the email address associated with their accounts. This configuration can be configured on a per-local identity profile basis.

About this task

Using the administrative console, configure the email ownership verifications settings for a local identity profile.

When you enable these settings, PingFederate generates a notification message for email ownership verification as the user submits the registration request. The email-verification message is valid for a configurable amount of time, 24 hours by default. If the user cannot find the previously sent message, the user can request another one by accessing the email ownership verification endpoint. Moreover, if profile management is enabled, the profile management page displays a reminder until the user verifies the associated email address as well. Like other local identity fields, the email verification status is stored in the directory and can be relayed to the applicable target applications through identity provider (IdP) authentication policies.

Steps

1. Go to Authentication # Policies # Local Identity Profiles.

2. On the Email Verification tab, select the Enable Email Ownership Verification check box if you want to offer users the opportunity to verify the email address associated with their accounts. This check box is not selected by default.



Note:

The rest of the steps apply only if you enable email ownership verification.

3. From the **Email Address Field** list, select a field.

The field value represents the recipient of the verification message.

Only fields that use the **Email** or **Text** input control are eligible and shown.

4. From the Ownership Status Field list, select a field

The field value represents the email ownership verification status. PingFederate sets the value to false in the directory when it receives a new or an updated email address from the user. Once the user verifies the email ownership, PingFederate sets the value to true.

Only fields that use the **Hidden** input control are eligible and shown.

5. If you want to modify the longevity of the link in the email-verification message, update the **One-Time Link Lifetime** field.

The default value is 1440 in minutes, 24 hours.

6. Optional: If you want to use different template files for various events, update the applicable template fields.

These templates are only applicable when using an SMTP Notification Publisher instance to deliver email-verification messages.

Default template fields and values are in the following table.

Email verification template fields and default values

Template field	Default value	
Email Template	message-template-email-ownership-verification.html	
Sent Template	local.identity.email.verification.sent.html	
Success Template	local.identity.email.verification.success.html	
Error Template	local.identity.email.verification.error.html	



Note:

You can find the email template file in the <pf install>/pingfederate/server/default/ conf/template/mail-notifications directory and the rest in the template directory.

7. From the **Notification Publisher** list, select an instance.

If you have not yet configured the desired notification publisher instance, click Manage Notification Publishers.

8. Click Done.

Configuring registration options

Using the registration settings, configure the user registration experience and specify the template file for the registration page.

Steps

Go to **Authentication # Policies # Local Identity Profiles**. On the **Registration** tab, configure the user registration page settings.

Choose from:

• If you want to enable invisible reCAPTCHA from Google to prevent automated registration attempts, select the **CAPTCHA** check box, and then click **Manage CAPTCHA Settings**.

This check box is not selected by default.

If you want to use a different template file, update the Registration Template field.

The default value is local.identity.registration.html.

Configuring profile management options

Configure the profile management experience and specify the template file for the profile management page.

About this task

Using the administrative console, configure the user registration in the local identity profiles section.

Steps

Go to **Authentication** # **Policies** # **Local Identity Profiles**. In the **Local Identity Profile** window, using the tabs, configure the user registration page settings. Choose from:

 If you want to give users the option to delete their local accounts without administrator assistance, select the Enable Profile Deletion check box.

This check box is not selected by default.

If enabled, when users choose to delete their accounts, their user records are removed from your directory.

If you want to use a different template file, update the Registration Template field.

The default value is local.identity.profile.html.

Managing datastore configuration

PingFederate requires datastore configuration because it stores customer identities in PingDirectory.

To begin, click Configure Data Store.

Selecting a datastore for customer identities

PingFederate stores customer identities in PingDirectory, improving the flexibility, scale, and security of your user data. Follow these steps to connect PingFederate to a directory to store identity and profile data and expose the data to all applications and channels through LDAPv3.

About this task

Using the Data Store Configuration tab, connect your LDAP datastore to PingFederate.

Steps

- 1. Go to Authentication # Policies # Local Identity Profiles.
- 2. On the Data Store Configuration tab, click Configure Data Store.

Result: This will open a Data Store window.

- 3. On the **Data Store Type** tab, enter a name in the **Name** field, and from the **Type** list, select **Directory** (LDAP).
- 4. Click Next.

If you have not yet created an LDAP datastore to connect PingFederate to your PingDirectory or if you want to review your LDAP datastore settings, click Manage Data Store.

Configuring LDAP base DN and attributes

Configure the datastore to search for a user's authentication starting with the base distinguished name (DN) and attributes within the LDAP directory.

About this task

On the LDAP Configuration tab, specify the branch of your directory hierarchy where you want PingFederate to store customer identities. Then, select the object class and the attributes to be associated with local identity fields.



Later you will associate the local identity profile with an HTML Form Adapter instance and apply the profile in an identity provider (IdP) authentication policy as part of the customer IAM configuration. If your use case requires registration or profile management, the policy engine must look up the users as they access the registration page or the profile management page. The scope of this search begins at the base DN defined here.

For this reason, the base DN here should match the value of the Search Base field defined in the LDAP Username Password Credential Validator instance used by the associated HTML Form Adapter instance.

For more information about each field, refer to the following table.

Field	Description
Base DN	The base distinguished name of the tree structure where PingFederate stores customer identities.
Root Object Class	The object class containing the desired attributes.
Attributes	A list of attributes based on the selected Root Object Class value.

Steps

- 1. Go to Authentication # Policies # Local Identity Profiles.
- 2. On the Data Store Configuration tab, click Configure Data Store.

Result: This will open a **Data Store** window.

- 3. On the **LDAP Configuration** tab, enter the applicable fields.
- 4. In the **User DN** field, specify a base DN.
- 5. Optional: Click View Local Identity Fields to determine which attributes from the directory server should be added to the local identity profile.
- 6. From the LDAP Configuration tab, click Advanced.

Result: This will open the LDAP Binary Attributes tab.

- 7. On the LDAP Binary Attributes tab, add attributes.
 - a. Enter a name in the Binary Attribute Name.
 - b. Click Add.
 - c. Select a root object class, select an applicable attribute, and then click Add Attribute.

Repeat this step to add more attributes as needed.

8. Click Done. Click Save.

Configuring LDAP relative DN and object class

When a user submits a registration request, PingFederate formulates the distinguished name (DN) of the user by prefixing the relative distinguished name (RDN) to the base DN defined in the LDAP configuration and then asks PingDirectory to create a new account based on the selected object class.

About this task

On the LDAP Configuration tab, enter a valid RDN pattern and select an object class.

Steps

- 1. Optional: Click View List of Available LDAP Attributes to determine which LDAP attributes can be used to construct the RDN pattern.
- 2. Enter a valid RDN pattern.

The pattern is as follows.

```
attribute1=value1[, ..., attributeN=valueN]
```

If you want to use the \${entryUUID} variable to guarantee the uniqueness of the relative DNs for all users, you must use it with the entryuuid LDAP attribute, such as in the following example.

```
entryUUID=${entryUUID}
```

- 3. Select an object class from the list.
- 4. Click Next.

Defining datastore mapping configuration

Configure the mapping between the local identity profile fields and the datastore attributes.

About this task

Steps

In the Data Store Configuration, on the Data Store Mapping tab, select an LDAP attribute under Data Store Attribute for each local identity field.

Reviewing datastore configuration

Review your datastore configuration to amend, keep, or discard your settings as needed.

About this task

On the **Summary** tab, review your data store configuration settings and then save them.

Steps

- 1. In the Data Store Configuration window, on the Summary tab, review your changes. Choose from:
 - To amend your configuration, click the corresponding tab title, then follow the configuration wizard to complete the task.
 - To keep your changes, click **Done** and continue with the rest of the configuration.



When editing an existing configuration, you can also click Save as soon as the administrative console offers the opportunity to do so.

- To discard your changes, click Cancel.
- 2. Click Done.

Review or make changes to your local identity profile configuration.

About this task

Using the administrative console, on the **Summary** tab, amend, save, or discard your local identity profile changes.

Steps

- 1. In the Data Store Configuration window, on the Summary tab, review your changes.
 - a. To amend your configuration, click the corresponding tab title, then follow the configuration wizard to complete the task.
 - b. To keep your changes, click **Done** and continue with the rest of the configuration.
 - c. To discard your changes, click Cancel.
- 2. When finished, click Done.

Configuring the HTML Form Adapter for customer identities

After defining a local identity profile, associate it with an instance of the HTML Form Adapter for PingFederate to leverage the HTML Form Adapter to present users the options to authenticate through third-party identity providers, self-register as part of the sign-on experience, and manage their accounts through a self-service profile management page.

About this task

Using the administrative console, on the IdP Adapter tab, create or add an HTML Form Adapter instance.



Note:

For registration and profile management, ensure the HTML Form Adapter instance is configured to validate credentials stored in PinaDirectory. This validation configuration is not required if your use case does not involve registration or profile management. For more information, see *Enabling third-party identity* providers without registration on page 352.

Steps

- 1. Go to the Authentication # Integration # IdP Adapters.
- 2. To create a new HTML Form Adapter instance, click **Create New Instance** or to reuse an existing instance, click on its name.
- 3. On the IdP Adapter tab, from the Password Credential Validator Instance list, select the LDAP Username Password Credential Validator instance that has been set up to validate credentials stored on your PingDirectory.



Note:

Skip this step if your use case does not involve registration or profile management.

- 4. On the IdP Adapter tab, from the Local Identity Profile list, select a local identity profile.
- 5. Click **Next** and complete the rest of the configuration tabs.
- 6. On the Summary tab, click Done. This will open the Manager IdP Adapter Instances window.
- 7. In the **Manager IdP Adapter Instances** window, click **Save** to save all changes.

Setting up self-service registration

PingFederate leverages the HTML Form Adapter to deliver a secure and easy-to-use customer authentication, registration, and profile management solution.

About this task

A typical self-service registration setup involves five components:

- A PingDirectory installation (step 1)
- An authentication policy contract (step 2)
- A local identity profile (step 3)
- An HTML Form Adapter instance (step 4)
- An IdP authentication policy (step 5)

To illustrate the configuration steps, consider the following example.

You need to support a consumer registration use case, where users complete a self-service registration process to create their accounts and then access resources protected by multiple service providers. For a registration to complete successfully, a user must provide an email address, a first name, a last name, an optional mobile phone number, and a password. The email address is the user identifier. All attributes are sent to the service providers as per the partner agreements. You have already created a specific object class in the directory to store the user information. The object class name is aPerson, and the LDAP attributes are mail, givenName, sn, and mobile.

Steps

- 1. Install PingDirectory.
- 2. Create an authentication policy contract.
 - a. Go to Authentication # Policies # Policy Contracts.
 - b. In the Policy Contracts window, click Create New Contract.
 - c. On the **Contract Attributes** tab, in the **Extend the Contract** field, extend the authentication policy contract with three additional attributes, such as, **firstName**, **lastName**, and **mobileNumber**.
 - d. After each entry add, click Add. Click Next.
 - e. On the Summary tab, review your changes. Click Done.
 - f. In the **Policy Contracts** window, click **Save**.

For more information, see .

- 3. Create a local identity profile using the **Authentication** # **Policies** # **Local Identity Profiles** configuration wizard.
 - a. In the Local Identity Profiles window, click Create New Profile.
 - b. On the **Profile Info** tab, enter a name in the **Local Identity Profile Name** field.
 - c. From the **Authentication Policy Contract** list, select the authentication policy (from *step 2*), and select the **Enable Registration** check box. Click **Next**.
 - d. On the Authentication Sources tab. click Next.
 - e. On the Fields tab, click Create New Field.
 - f. In the **Field Configuration** window, on the **Field Configuration** tab, define four local identity fields. Enter the information described in the following table.

Local Identity Profile fields and entries

Туре	ID	Label	Parameters
Email	lipEmail	Email address	Select the Required check box.
Text	lipFirstName	First name	Select the Required check box.
Text	lipLastName	Last name	Select the Required check box.

Туре	ID	Label	Parameters
Phone	lipMobile	Mobile number	No parameters are required.

- g. After each field entry, click **Next**. On the **Summary** tab, review your changes. Click **Done**.
- h. Repeat steps e through g until the fields are entered.

As needed, select the **Mask Log Values** check box for any of the four local identity fields and the **Mask all OGNL-expression generated log values** check box. The latter applies to all local identity fields.

- i. In the **Local Identity Profile** window, on the **Fields** tab, identify an ID field to be the unique ID for your configuration, and click the corresponding **Unique ID**. Click **Next**.
- j. On the **Email Verification** tab, click **Next**.
- k. On the Registration tab, click Next.
- I. On the **Data Store Configuration** tab, click **Configure Data Store**.
- m. In the **Data Store Configuration** window, on the **Data Store** tab, from the **Data Store** list, select the LDAP datastore that has been set up to connect to your PingDirectory. Click **Next**.
- n. In the **Data Store Configuration** window, on the **LDAP Configuration** tab, specify the branch of your directory hierarchy where you want PingFederate to store customer identities in the **Base DN** field and the LDAP attributes to be associated with fields defined in this local identity profile under **Attribute**.
- o. In the **Data Store Configuration** window, on the **Identity Creation** tab, define the RDN pattern in the **Relative DN Pattern** field, and select your object, such as class aPerson for this sample use case, from the **Object Class** list.

The pattern is as follows.

```
attribute1=value1[, ..., attributeN=valueN]
```

If you want to use the \${entryUUID} variable to guarantee the uniqueness of the relative DNs for all users, you must use it with the entryUUID LDAP attribute.

```
entryUUID=${entryUUID}
```

p. In the **Data Store Configuration** window, on the **Data Store Mapping** tab, configure the mapping between the local identity profile fields and the datastore attributes. See the following table.

Mapping entries for local identity profile fields and datastore attributes

Field	Data Store Attribute
lipEmail	mail
lipFirstName	givenName
lipLastName	sn
lipMobile	mobile

- q. In the Data Store Configuration window, on the Summary tab, click Done.
- r. On the **Summary** tab of the local identity profile, click **Save**.

For more information, see *Defining a local identity profile* on page 327.

- 4. Configure an HTML Form Adapter instance for customer identities.
 - a. Go to the IdP Adapters window.
 - b. Create a new HTML Form Adapter instance or reuse an existing one by clicking its name.
 - c. On the IdP Adapter tab, add the LDAP Username Password Credential Validator instance that has been set up to validate credentials stored on your PingDirectory.
 - d. On the IdP Adapter tab, select the newly created local identity profile from the Local Identity Profile list.
 - e. Complete the rest of the configuration and save all changes.

(For more information, see .)

- 5. Create an IdP authentication policy.
 - a. Go to Authentication # Policies # Policies.
 - b. Select the HTML Form Adapter instance (configured in step 4) under Action.
 - 1. For its **Fail** path, select **Done**.
 - 2. For its **Success** path, select the local identity profile (created in step 3).
 - c. Click **Local Identity Mapping** underneath the selected local identity profile, which opens the **Inbound Mapping & Contract Fulfillment** configuration wizard.
 - d. On the **Inbound Mapping & Contract Fulfillment Inbound Mapping** window, configure the pf.local.identity.unique.id built-in local identity field for the registration process.

At runtime, PingFederate fulfills the value of the pf.local.identity.unique.id built-in local identity field based on this configuration and passes the value to PingDirectory. PingDirectory uses this value to determine whether such identity has already been created. The pf.local.identity.unique.id field value should therefore be mapped from the subject identifier of the preceding authentication source, namely the username attribute from the HTML Form Adapter.

For this sample use case, configure the **Inbound Mapping** window as in the following table.

Inbound Mapping Fulfillment	Source	Value
pf.local.identity.unique.id	Adapter	username

- e. On the Attribute Sources & User Lookup tab, click Next.
- f. On the **Contract Fulfillment** tab, fulfill the authentication policy contract with values from this local identity profile as follows:

Outbound Contract Fulfillment	Source	Value
subject	Local Identity	lipEmail
firstName	Local Identity	lipFirstName
lastName	Local Identity	lipLastName

Outbound Contract Fulfillment	Source	Value
mobileNumber	Local Identity	lipMobile

- g. On the Issuance Criteria tab, click Next.
- h. On the Summary tab, click Done.

Result: The Inbound Mapping & Contract Fulfillment configuration wizard brings back the Manage Authentication Policies window.

i. Select the IdP Authentication Policies check box.



Other IdP authentication policies, if any, are enabled as well.

j. Click **Save** to keep your changes.

For more information, see Applying policy contracts or identity profiles to authentication policies on page 250.

6. Map the authentication policy contract to the applicable Browser SSO connections, OAuth grantmapping configuration, or both. See Managing authentication source mappings on page 417 and Managing authentication policy contract grant mapping on page 538.

Result

You have now successfully set up self-service registration. When users sign on through this HTML Form Adapter instance, they have the option to complete a self-service registration process to create their accounts using the Register now link, as illustrated in the following screen capture.



If a user chooses to register, the HTML Form Adapter redirects the user to the registration page. Based on the configuration of this sample use case, as illustrated in the following registration screen capture.



For registration, you can optionally allow users to leverage their existing identities from third-party identity providers. Any identity provider (IdP) connection or IdP adapter, such as the LinkedIn Cloud Identity Connector, can be used as an authentication source to a third-party identity provider.

About this task

Using the Administrative console, enable third-party IdP. This configuration involves the same five components to set up registration, plus the IdP connections or IdP adapter instances to connect with the third-party identity providers. This capability enables a mapping configuration between the attributes returned by the identity provider and the fields within the registration page, streamlining the registration process. See the following configuration steps.

- IdP connections or IdP adapter instances
- A PingDirectory installation (step 1)
- An authentication policy contract (step 2)
- A local identity profile (step 3)
- An HTML Form Adapter instance (step 4)
- An IdP authentication policy (step 5)

To illustrate the configuration steps, consider the following example.

You need to support a consumer registration use case, where users complete a self-service registration process to create their accounts and then access resources protected by multiple service providers. For a registration to complete successfully, a user must provide an email address, a first name, a last name, an optional mobile phone number, and a password. The email address is the user identifier. All attributes are sent to the service providers as per the partner agreements. You have already created a specific object class in the directory to store the user information. The object class name is aPerson, and the LDAP attributes are mail, givenName, sn, and mobile.

Additionally, this use case must also allow users to take advantage of their existing accounts at ACME, a major social network, for registration and authentication. It happens that you have already established an IdP connection to this social network, from which you received the same set of attributes: SAML SUBJECT for the user's email address, ssoFirstName, ssoLastName, and ssoMobile.

Configure a third-party identifier.



If you know how to set up PingDirectory to connect with PingFederate and an authentication policy contract shown in Setting up self-service registration on page 337, you can skip to step 3 to create a local identity profile.

Steps

- 1. Install PingDirectory.
- 2. Create an authentication policy contract.
 - a. Go to Authentication # Policies # Policy Contracts.
 - b. In the Policy Contracts window, click Create New Contract.
 - c. On the Contract Attributes tab, in the Extend the Contract field, extend the authentication policy contract with three additional attributes, such as, firstName, lastName, and mobileNumber.
 - d. After each entry add, click Add. Click Next.
 - e. On the **Summary** tab, review your changes. Click **Done**.
 - f. In the **Policy Contracts** window, click **Save**.

For more information, see .

- 3. Create a local identity profile using the Authentication # Policies # Local Identity Profiles. Click Create New Profileto access the Local Identity Profile window and configuration wizard.
 - a. On the Profile Info tab, enter a name in the Local Identity Profile Name field, select the authentication policy from step 2, and select the Enable Registration check box. Click Next.
 - b. On the Authentication Sources tab, enter ACME in the Authentication Source field. Click Add.



To support additional third-party identity providers, enter a value for each. At runtime, the sign-on page displays them in the order defined on this window.



Tip:

If you know how to set up a local identity profile and an HTML Form Adapter instance for customer identities shown in Setting up self-service registration on page 337), you can skip to step 5 to create an IdP authentication policy.

c. On the **Fields** tab, define four local identity fields as shown in the following table.

Туре	ID	Label	Parameters
Email	lipEmail	Email address	Select the Required and Unique ID check boxes.
			Note: The Unique IDcheck boxes are on the Fields tab.
Text	lipFirstName	First name	Select the Required check box.
Text	lipLastName	Last name	Select the Required check box.
Phone	lipMobile	Mobile number	None required.

As needed, select the Mask Log Values check box for any of the four local identity fields, and select the Mask all OGNL-expression generated log values check box for all fields.

- d. On the Email Verification tab. click Next.
- e. On the **Registration** tab, click **Next**.
- f. On the Data Store Configuration tab, click Configure Data Store.
- g. In the Data Store Configuration window, on the Data Store tab, from the Data Store list, select the LDAP datastore that has been set up to connect to your PingDirectory. Click Next.
- h. In the Data Store Configuration window, on the LDAP Configuration tab, specify the branch of your directory hierarchy where you want PingFederate to store customer identities in the Base DN

field and the LDAP attributes to be associated with fields defined in this local identity profile under **Attribute**.

i. In the **Data Store Configuration** window, on the **Identity Creation** tab, define the RDN pattern in the **Relative DN Pattern** field, and select your object, such as class aPerson for this sample use case, from the **Object Class** list.

The pattern is as follows.

```
attribute1=value1[, ..., attributeN=valueN]
```

If you want to use the \${entryUUID} variable to guarantee the uniqueness of the relative DNs for all users, you must use it with the entryUUID LDAP attribute.

```
entryUUID=${entryUUID}
```

j. In the **Data Store Configuration** window, on the **Data Store Mapping** tab, configure the mapping between the local identity profile fields and the datastore attributes. See the following table.

Mapping entries for local identity profile fields and datastore attributes

Field	Data Store Attribute
lipEmail	mail
lipFirstName	givenName
lipLastName	sn
lipMobile	mobile

- k. In the Data Store Configuration window, on the Summary tab, click Done.
- I. On the **Summary** tab of the local identity profile, click **Save**.

For more information, see *Defining a local identity profile* on page 327.

- 4. Configure an HTML Form Adapter instance for customer identities.
 - a. Go to the IdP Adapters window.
 - b. Create a new HTML Form Adapter instance or reuse an existing one by clicking its name.
 - c. On the IdP Adapter tab, add the LDAP Username Password Credential Validator instance that has been set up to validate credentials stored on your PingDirectory.
 - d. On the IdP Adapter tab, select the newly created local identity profile from the Local Identity Profile
 - e. Complete the rest of the configuration and save all changes.

(For more information, see .)

- 5. Create an IdP authentication policy.
 - a. Go to the Authentication # Policies # Policies.
 - b. Under **Policy**, select the HTML Form Adapter instance configured in step 4.
 - c. In the Policy window, under the Value section, select the drop-down arrow to show the Fail and Success fields.
 - 1. For the **Fail** path, select **Done**.
 - 2. For the **Success** path, select the local identity profile created in *step 3*. Click **Done**.
 - d. Click Local Identity Mapping underneath the selected local identity profile, which opens the Inbound Mapping & Contract Fulfillment configuration wizard.



Note:

The next few steps configure the fulfillment of the authentication policy contract for the scenario where users choose to register directly without going through ACME.



If you know how to setup the inbound mapping and contract fulfillment of an authentication policy contract through a local identity profile shown in Setting up self-service registration on page 337, you can skip to step i to create a rule for the scenario where users choose to register and subsequently authenticate via ACME.

e. On the Inbound Mapping & Contract Fulfillment Inbound Mapping window, configure the pf.local.identity.unique.id built-in local identity field for the registration process. At runtime, PingFederate fulfills the value of the pf.local.identity.unique.id builtin local identity field based on this configuration and passes the value to PingDirectory. PingDirectory uses this value to determine whether such identity has already been created. The pf.local.identity.unique.id field value should therefore be mapped from the subject

identifier of the preceding authentication source, namely the username attribute from the HTML Form Adapter.

For this sample use case, configure the **Inbound Mapping** window as in the following table.

Inbound Mapping Fulfillment	Source	Value
pf.local.identity.unique.id	Adapter	username

- f. On the Attribute Sources & User Lookup tab, click Next.
- g. On the Attribute Sources & User Lookup tab, click Next.
- h. On the Issuance Criteria tab, click Next.
- i. On the Inbound Mapping & Contract Fulfillment # Summary window, click Done.

The Inbound Mapping & Contract Fulfillment configuration wizard brings back the Manage Authentication Policies window.



Note:

The remaining steps configure the fulfillment of the authentication policy contract for the scenario where users choose to register and subsequently authenticate through ACME.

- j. Click **Rules** underneath the **Success** path of the HTML Form Adapter instance.
- k. On the Rules dialog, create a policy path for users who choose to register and authenticate via ACME. For this sample use case, configure as in the following table.

Authentication policy rules fields and entries

Attribute Name	Condition	Value	Result
policy.action	equal to	ACME	ACME users
		Important: The value here must match the value defined on the Authentication Sources window. See step 3b.	The Result field controls the label shown for the policy path of this rule. The value does not need to match the value defined on the Authentication Sources window.



Important:

If you have defined multiple third-party identity providers on the Authentication Sources window, you must repeat these steps to add a policy.action rule to create a policy path for each.

In addition, select the **Default to Success** check box, the default behavior. When selected, the Success path remains, which is important for this sample use case where users are free to choose whether to register and subsequently authenticate via ACME.

When finished, click **Done**, which brings you back to the **Authentication Policies** window.

I. For the **ACME** users path, select the IdP connection to ACME under **Action**.



Z Tip:

Generally speaking, any IdP adapter instance or IdP connection that connects to the third-party identity provider can be used here.



Note:

If you have defined multiple third-party identity providers and added rules to create a policy path for each, you may select Restart. The Restart policy action provides users the opportunity to do over. When triggered, the policy engine routes the requests back to the first checkpoint of the invoked authentication policy.

By selecting Restart for the Fail path, you give users the opportunity to choose another thirdparty identity provider when they fail to authenticate through ACME.

Undesirable looping behaviors can occur if you select **Restart** for the **Fail** path at the root of an authentication policy tree. PingFederate mitigates this risk by automatically limiting the number of policy restarts per transaction.

- 2. For its **Success** path, select the local identity profile created in *step 3*.
- m. Click Local Identity Mapping underneath the selected IdP connection, which opens the Inbound Mapping & Contract Fulfillment configuration wizard.
- n. On the Inbound Mapping & Contract FulfillmentInbound Mapping window, configure the pf.local.identity.unique.id built-in local identity field for the registration process and optionally other fields so that PingFederate can pre-populate values for these fields on the registration page.

At runtime, PingFederate fulfills the value of the pf.local.identity.unique.id builtin local identity field based on this configuration and passes the value to PingDirectory. PingDirectory uses this value to determine whether such identity has already been created. The pf.local.identity.unique.id field value should therefore be mapped from the subject identifier of the preceding authentication source, namely the subject identifier from the IdP connection.

For this sample use case, the **Inbound Mapping** window is configured as in the following table.

Inbound Mapping fields and entries

Inbound Mapping Fulfillment	Source	Value
pf.local.identity.unique.id	IdP Connection	SAML_SUBJECT
lipEmail	IdP Connection	SAML_SUBJECT
lipFirstName	IdP Connection	ssoFirstName
lipLastName	IdP Connection	ssoLastName

Inbound Mapping Fulfillment	Source	Value
lipMobile	IdP Connection	ssoMobile

- o. On the Attribute Sources & User Lookup tab, click Next.
- p. On the Attribute Sources & User Lookup tab, click Next.
- q. On the Issuance Criteria tab, click Next.
- r. On the Inbound Mapping & Contract Fulfillment Summary tab, click Done.

The Inbound Mapping & Contract Fulfillment configuration wizard brings back the Manage Authentication Policies window.



Important:

If you have defined multiple rules, each forming a policy path for a third-party identity provider, ensure you complete the Inbound Mapping & Contract Fulfillment configuration for each of them.

s. In the Policies window, on the Policies tab, select the IdP Authentication Policies check box.



Note:

Other IdP authentication policies, if any, are enabled as well.

t. Click Save to keep your changes.

For more information, see Applying policy contracts or identity profiles to authentication policies on page 250.

Result

You have now successfully set up self-service registration with an option for users to register and subsequently authenticate via ACME. When users sign on through this HTML Form Adapter instance, they have two registration options:

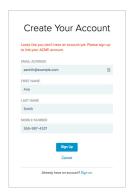
- Click the Register now link, fill in the registration page, and register.
- Click the social sign-on link, authenticate via ACME, review the registration page, and register.

Based on the configuration of this sample use case, the following screen capture of a sign-on page is presented.



If you have added Facebook, Google, LinkedIn, and Twitter as the authentication sources, the following sign-on page is presented.

Suppose a user chooses to register through ACME. Once authenticated and redirected back to PingFederate, PingFederate pre-populates the registration page with values it receives from ACME, as illustrated in this screen capture.



This registration option streamlines the self-service registration process.

Enabling profile management

In addition to registration, you can enable self-service profile management and specify which local identity fields users can update on the profile management page.

About this task

Use the administrative console to enable profile management.

To illustrate the configuration steps, consider the sample use case in Setting up self-service registration on page 337 or Enabling third-party identity providers on page 341 with the added requirement of allowing users to modify their mobile number and to remove their local accounts.

Configuration steps:



As the required components remain the same, the step sequence matches those in Setting up self-service registration on page 337 and Enabling third-party identity providers on page 341 as well. If you require more information for a given step, see the same step in one of the aforementioned pages.

Steps

- 1. Set up PingDirectory to connect with PingFederate.
- 2. Create an authentication policy contract. For more information on how to create an authentication policy contract, see Managing policy contracts.

- 3. Configure profile management when creating a new or reconfiguring an existing local identity profile.
 - a. Go to Authentication # Policies # Local Identity Profiles configuration wizard.
 - b. On the Profile Info tab, select the Enable Profile Management check box. Click Next.
 - c. Optional: On the **Authentication Sources** tab, define authentication sources. For more information, see *Defining authentication sources*.
 - d. On the **Fields** tab, select the **Profile Management** check box under **Applies To** for the applicable fields as you define local identity fields.

These selected local identity fields will be shown to authenticated users on the profile management page.

For this sample use case, select the **Profile Management** check box for the lipMobile local identity field.

- e. On the Email Verification tab, click Next.
- f. On the **Registration** tab, click **Next**.
- g. On the **Profile Management** tab, select the **Enable Profile Deletion** check box.

In general, this is an optional feature. It is selected here because it is one of the requirements of this sample use case.

- h. Continue from step 3f as documented in *Setting up self-service registration* on page 337 or *Enabling third-party identity providers* on page 341.
- 4. Configure an HTML Form Adapter instance for customer identities. For more information, see *Configuring the HTML Form Adapter for customer identities*.
- 5. Create an IdP authentication policy. For more information, see *Defining authentication policies*.
- 6. Provide the profile management URL to users.
 - a. Go to Authentication # Policies # Local Identity Profiles.
 - b. Select the local identity profile that you have configured profile management in step 3.
 - c. Copy the profile management URL as shown on the Summary tab and pass it to the users.

Result

You have now successfully enabled profile management. Authenticated users can review and modify the local identity fields that have been configured to show on the profile management page and delete their local accounts if the option has been enabled.

The following screen capture provides a sample of the profile management page based on the sample use case.

If you add Facebook, Google, LinkedIn, and Twitter to the local identity profile, when a user accesses the profile management page, the user will see a page similar to the following screen capture.

If you have only one authentication source, the profile management page reminds the users that they must set a password for their local accounts before disconnecting the third-party identity provider.

Creating advanced registration mapping

PingFederate leverages the HTML Form Adapter to deliver a secure and easy-to-use customer authentication, registration, and profile management solution. The HTML Form Adapter contract includes two core attributes.

About this task

Using the administrative console, create advanced registration mapping by defining authentication policies in your adapter instances.

To illustrate the configuration steps, consider the following setup that you have already made with the parameters username and policy.action. Whether or not the local identity profile is configured with any authentication sources, if the user chooses to register directly by clicking on the Register now link, PingFederate sets the value to identity.registration. This fulfillment allows you to create rules to differentiate authentication requirements from the registration flow.

- A PingDirectory installation with a set of users.
- An LDAP datastore, an LDAP Username Password Credential Validator instance, and an HTML Form Adapter instance on PingFederate to validate credentials stored in PingDirectory.
- An IdP authentication policy that chains the HTML Form Adapter instance, an PingID Adapter instance, and an authentication policy contract for the purpose of enforcing PingID multi-factor authentication in multiple browser-based single sign-on (SSO) use cases through service provider (SP) connections, OAuth authorization code flow, and OAuth implicit flow. The following window capture illustrates your existing policy.



To illustrate the configuration steps, consider the following setup that you have. You need to add support for a consumer registration use case similar to the one in Setting up self-service registration on page 337, and, at the same time, keep the policy that enforces the multi-factor authentication requirement.

Configuration steps.

Steps

- 1. Set up PingDirectory for customer identities.
- 2. Make a note of which authentication policy contract is currently being used in your policy.
- 3. Create a local identity profile.
 - a. G to Authentication # Policies # Local Identity Profiles configuration wizard.
 - b. On the Profile Info tab, in the Local Identity Profile Name field, enter a name of the local identity profile, and from the Authentication Policy Contract list, select the authentication policy from step 2. Select the Enable Registration check box.

If you want to enable profile management as well, select the relevant check box.

- 4. Configure the HTML Form Adapter instance for customer identities.
 - a. Go to Authentication # Integration # IdP Adapters.
 - b. From the **Instance Name** section, select the **HTMLFormAdapter**.
 - c. On the IdP Adapter tab, from the Local Identity Profile list, select a local identity profile.
 - d. Complete the rest of the configuration and save all changes.

- 5. Modify your existing IdP authentication policy.
 - a. Go to Authentication # Policies # Policies.
 - b. On the Policies tab, under the Policy section, click the relevant policy to open the Policy window.
 - c. Under the **Success** path of the HTML Form Adapter instance, click **Rules**.
 - d. In the **Rules** dialog, create a policy path for users who choose to register. For this sample use case, configure as follows.

Defining authentication policy rules attributes field and values

Attribute Name	Condition	Value	Result
policy.action	equal to	identity.registration	Registration
			Complete the rest of the configuration to create the local identity. The Result field controls the label shown for the policy path of this rule.

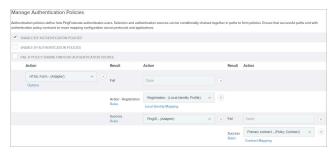
In addition, ensure the **Default to Success** check box is selected. When selected, the **Success** path remains, which is important for this sample use case where users are redirected to the

PingID Adapter instance to fulfill the multi-factor authentication requirement after authenticating successfully against the HTML Form Adapter.

When finished, click **Done**, which brings you back to the **Manage Authentication Policies** window.

e. For the **Registration** path, select the local identity profile from step 3 under **Action** and then complete its **Local Identity Mapping**configuration.

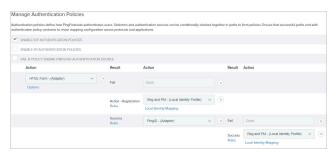
The following screen capture illustrates your new policy.



f. If you have enabled profile management in step 3, you must also replace the policy contract with the local identity profile and then complete its **Local Identity Mapping** configuration.

This step is required so that PingFederate can route users through the HTML Form # PingID policy path when they try to access the profile management page.

The following screen capture illustrates this change.





Note:

No reconfiguration is required in your Browser SSO connections and OAuth grant-mapping configuration for your new policy to take effect.

g. Click Save to keep your changes.

Result

You have now successfully added the requested consumer registration and profile management use case to your current policy.

Enabling third-party identity providers without registration

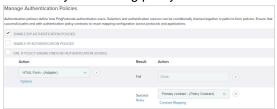
If you have already configured identity provider (IdP) connections or IdP adapters to connect with thirdparty identity providers, you can enhance the HTML Form Adapter sign-on page with the option to authenticate with these providers.

Before you begin

Consider the following setup that you have already made.

An HTML Form Adapter instance to validate local user credentials.

- An authentication policy contract.
- An IdP authentication policy that chains the HTML Form Adapter instance and an authentication policy contract so that the policy contract can harness attribute values returned by the HTML Form Adapter instance for multiple browser-based single sign-on (SSO) use cases via service provider (SP) connections, OAuth authorization code flow, and OAuth implicit flow. The following window capture illustrates your existing policy.



About this task

This setup involves the following components:

- IdP connections or IdP adapter instances configured to connect with your third-party identity providers
- An authentication policy contract
- A local identity profile
- An HTML Form Adapter instance
- An IdP authentication policy

You need to enhance the sign-on experience by giving users the option to authenticate using their existing accounts at ACME, a major social network. It happens that you have already established an IdP connection to this social network.

Configuration steps:

Steps

1. Verify the IdP connection returns the attributes required to complete the browser-based SSO use cases.



As needed, you can also deploy and configure Cloud Identity Connectors to support identities from Facebook, Google, LinkedIn, or Twitter.

- 2. Make a note of which authentication policy contract is currently being used in your policy.
- 3. Create a local identity profile using the Authentication # Policies # Local Identity Profiles configuration wizard. Click Create New Profile.
 - a. On the Profile Info tab, enter a name in the Local Identity Profile Name field and from the Authentication Policy Contract, select the authentication policy from step 2. Click Next.
 - b. On the Authentication Sources tab, enter ACME under Authentication Source, and then click **Add**. Click **Done** to exit the configuration window.



Note:

To support additional third-party identity providers, enter a value for each. At runtime, the sign-on page displays them in the order defined on this window.

- 4. Configure the HTML Form Adapter instance for customer identities.
 - a. Go to Integration # IdP Adapters.
 - b. On the IdP Adapters window, from the Instance Name list, click the HTMLFormAdapter instance.
 - c. On the IdP Adapter tab, from the Local Identity Profile list, select a local identity profile.
 - d. Complete the rest of the configuration and save all changes.
- 5. Modify your existing IdP authentication policy.
 - a. Go to Authentication # Policies # Policies.
 - b. On the **Policies** tab, under the **Policy** section, click the existing IdP policy.
 - c. Click **Rules** under the **Success** path of the HTML Form Adapter instance.
 - d. In the Rules dialog, create a policy path for users who choose to authenticate with ACME. For this sample use case, configure the fields as in the following table.

Defining authentication policy rules fields and entries

Attribute Name	Condition	Value	Result
policy.action	equal to	ACME	ACME users
		Important: The value here must match the value defined on the Authentication Sources window. See step 3b.	The Result field controls the label shown for the policy path of this rule. The value does not need to match the value defined on the Authentication Sources window.



Important:

If you have defined multiple third-party identity providers on the Authentication Sources tab, you must repeat these steps to add a policy.action rule to create a policy path for each.

In addition, ensure the **Default to Success** check box is selected. When selected, the **Success** path remains, which is important for this sample use case where users can also authenticate using their local accounts.

- e. When finished, click Done. This will bring you back to the Authentication Policies window.
- f. For the **ACME** users path, select the IdP connection to ACME under **Action**.



Generally speaking, any IdP adapter instance or IdP connection that connects to the third-party identity provider can be used here.

1. For its **Fail** path, select **Done**.



If you have defined multiple third-party identity providers and added rules to create a policy path for each, you can select Restart. The Restart policy action provides users the opportunity to do over. When triggered, the policy engine routes the requests back to the first checkpoint of the invoked authentication policy.

By selecting **Restart** for the **Fail** path, you give users the opportunity to choose another third-party identity provider when they fail to authenticate through ACME.

2. For its **Success** path, select the local identity profile created in *step 3* and then complete the **Local Identity Mapping** configuration.



Note:

Because this use case does not involve registration, the source of fulfillment is limited to the preceding IdP connection or IdP adapter instance, dynamic text, and attribute mapping expression, if enabled.

The following screen capture illustrates your new policy.

g. Click Save to keep your changes.

Result

You have now successfully added the option to authentication via ACME without enabling registration. When users sign on through this HTML Form Adapter instance, the following sign-on page is presented.



If you have added Facebook, Google, LinkedIn, and Twitter as the authentication sources, the following sign-on page is presented.



Users can sign on using their local accounts or third-party identity provider accounts.

Customizing assertions and authentication requests

Customize applicable messages by enabling OGNL expression and going to the **URL** window to access the **Show Advanced Customizations** option.

About this task

Some browser single sign-on (SSO) use cases might require additional customizations in the assertions sent from the PingFederate identity provider (IdP) server to the service provider (SP), or in the authentication requests sent from the PingFederate SP server to the IdP. PingFederate can fulfill these use cases on a per-connection basis using OGNL expressions.

- 1. Enable OGNL expression by editing the org.sourceid.common.ExpressionManager.xml file, located in the <pf install>/pingfederate/server/default/data/config-store directory.
- 2. Select the applicable SP or IdP connection.
- 3. On the Activation & Summary window, scroll to the Protocol Settings section, and click Assertion Consumer Service URL for an SP connection, or click SSO Service URLs for an IdP connection.
- 4. Click **Show Advanced Customizations** to customize the applicable message.

The available customizable Message Types vary depending on your federation role (IdP or SP) as well as the protocol of the connection (SAML 1.x, SAML 2.0, and WS-Federation). After you select a message type, you have access to its list of Available Variables. You can customize the assertions or the authentication requests to fulfill your use case,

Message types and available variables

Advanced customizations depend on available message types and available variables for both service provider (SP) connections and identity provider (IdP) connections.

The following tables describe the relationship between message type and available variable, as well as the corresponding class or interface information in Javadoc.



The Javadoc for PingFederate is located in the <pf install>/pingfederate/sdk/doc directory.

SP connections (SAML 2.0)

Message Types	Available Variables
	Classes/Interfaces in Javadoc
AssertionType	#AssertionType
	org.sourceid.saml20.xmlbinding.assertion.AssertionType
	#AssertionTypes
	org.sourceid.saml20.xmlbinding.assertion.AssertionType[]
	#Attributes
	org.sourceid.util.log.AttributeMap
ResponseDocument	#ResponseDocument
	org.sourceid.saml20.xmlbinding.protocol.ResponseDocument
	#Attributes
	org.sourceid.util.log.AttributeMap

SP connections (SAML 1.x)

Message Types	Available Variables
	Classes/Interfaces in Javadoc
AssertionType	#AssertionType
	org.sourceid.protocol.saml11.xml.AssertionType
	#AssertionTypes
	org.sourceid.protocol.saml11.xml.AssertionType[]
	#Attributes
	org.sourceid.util.log.AttributeMap
ResponseDocument	#ResponseDocument
	org.sourceid.protocol.samlp11.xml.ResponseDocument
	#Attributes
	org.sourceid.util.log.AttributeMap

SP connections (WS-Federation)

Message Types	Available Variables	
	Classes/Interfaces in Javadoc	
AssertionType	#AssertionType	
	org.sourceid.protocol.saml11.xml.AssertionType	
	#Attributes	
	org.sourceid.util.log.AttributeMap	
RequestSecurityToken ResponseDocument	#RequestSecurityTokenResponseDocument	
	org.xmlsoap.schemas.ws.x2005.x02.trust.RequestSecurityTokenRespons	eDocument
	#Attributes	
	org.sourceid.util.log.AttributeMap	

IdP connections (SAML 2.0)

Message Type	Available Variables
	Classes/Interfaces in Javadoc
AuthnRequestDocument#AuthnRequestDocument	
	org. source id. saml 20. xml binding. protocol. Authn Request Document

Other available variables (regardless of roles and protocols)

Available Variables	Classes/Interfaces in Javadoc
#XmlHelper	com.pingidentity.sdk.xml.XmlHelper
#HttpServletRequest	javax.servlet.http.HttpServletRequest

Available Variables	Classes/Interfaces in Javadoc
#HttpServletResponse	javax.servlet.http.HttpServletResponse

Variables related to Federation Hub (regardless of message type)

Connections	Protocol	Available Variables
		Classes/Interfaces in Javadoc
SP and IdP connections	SAML 2.0	#FedHubIncomingAuthnRequest
		org. source id. saml 20. xml binding. protocol. Authn Request Document and the protocol of t
SP connection	SAML 2.0	#FedHubOutgoingAuthnRequest
		org. source id. saml 20. xml binding. protocol. Authn Request Document and the protocol of t
SP connection	SAML 2.0	#FedHubIncomingAuthnResponse
	SAML 1.x WS-Federation	org.sourceid.saml20.xmlbinding.protocol.ResponseDocument (SAML 20)
		org.sourceid.protocol.samlp11.xml.ResponseDocument (SAML 1.x)
		org.xmlsoap.schemas.ws.x2005.x02.trust. RequestSecurityTokenResponseDocument (WS-Federation)
SP connection	SAML 2.0	#FedHubIdpConnPartnerId
	SAML 1.x	java.lang.String
	WS-Federation	The Partner's Entity ID in the IdP connection that bridges the identity provider.
SP connection	SAML 2.0	#FedHubIdpConnProtocol
	SAML 1.x	java.lang.String
	WS-Federation	The protocol of the SP connection. The returned values are SAML20, SAML11, SAML10, or WSFED.
IdP connection	SAML 2.0	#FedHubSpConnPartnerId
	SAML 1.x	java.lang.String
	WS-Federation	The Partner's Entity ID in the SP connection that bridges the service provider.
IdP connection	SAML 2.0	#FedHubSpConnProtocol
	SAML 1.x	java.lang.String
	WS-Federation	The protocol of the IdP connection. The returned values are SAML20, SAML11, SAML10, or WSFED.

Sample customizations

Use OGNL expressions to customize assertions and authentication requests in different ways.

Add SessionNotOnOrAfter to assertions

This expression adds the optional SessionNotOnOrAfter attribute to the <AuthnStatement> element and sets the value to 60 minutes.

Message Type

AssertionType

Expression

```
#cal = new org.apache.xmlbeans.XmlCalendar(new java.util.Date()),
#cal.setTimeZone(@java.util.TimeZone@getTimeZone("UTC")),
#cal.add(@java.util.Calendar@MINUTE, 60),
#AssertionType.getAuthnStatementArray(0).setSessionNotOnOrAfter(cal)
```

Expected assertions

Use well-formed XML as attribute value

The following expression inserts well-formed XML in the <attributeValue> element if the Attribute Name Format is urn:pingidentity.com:SAML:attrname-format:xml:complex.

Message Type

AssertionType

Expression

```
#i = 0,
#AssertionType.getAttributeStatementArray(0).getAttributeArray().{
    #this.getNameFormat().equals('urn:pingidentity.com:SAML:attrname-
format:xml:complex')?{
        #xml = #this.getAttributeValueArray(0).getStringValue(),
        #ast =
    @org.sourceid.saml20.xmlbinding.assertion.AttributeStatementType
$Factory@parse(#xml),
        #AssertionType.getAttributeStatementArray(0).setAttributeArray(#i,
        ast.getAttributeArray(0))
    }:null,
#i = #i+1
}
```



Note:

Line breaks are inserted for readability only. Statements calling methods whose arguments are enclosed in quotes must be entered on a single line.

This example uses well-formed XML as the attribute value for attributes that are configured as urn:pingidentity.com:SAML:attrname-format:xml:complex (a custom attribute name format added to $< pf_install > pingfederate/server/default/data/config-store/custom-name-formats.xml) in the Attribute Contract window. You can use other application logic here.$

Sample inputs (attributes and their values)

Attribute Name	ExtAttr1
Attribute Name Format	urn:oasis:names:tc:SAML:2.0:attrname-format:unspecified
Attribute Value	123

```
Attribute Name
                    ExtAttr2
Attribute Name Format urn:pingidentity.com:SAML:attrname-format:xml:complex
Attribute Value
                     <saml:Attribute</pre>
                      xmlns:saml="urn:oasis:names:tc:SAML:2.0:assertion"
                      Name="ExtAttr2"
                      NameFormat="urn:oasis:names:tc:SAML:2.0:attrname-
                     format:unspecified">
                         <saml:AttributeValue</pre>
                          xmlns:xsi="http://www.w3.org/2001/XMLSchema-
                     instance"
                          xmlns:customNs="http://www.sample.tld/
                     customnamespace">
                             <customNs:Line>Documentation/customNs:Line>
                              <customNs:Line>Ping Identity</customNs:Line>
                         </saml:AttributeValue>
                     </saml:Attribute>
                    Note:
                    This is a well-formed XML document in one line.
```

Expected results

```
</saml:AttributeValue>
</saml:Attribute>
```

Include extensions in authentication requests

This expression includes the optional Extensions element in the authentication requests if a certain query parameter (oid in this example) is sent to the /sp/startSSO.ping endpoint to start an SPinitiated SSO request.

Message Type

AuthnRequestDocument

Expression

```
#element = #XmlHelper.addToSaml2Extensions(#AuthnRequestDocument,
'<samplens:orgId name="orgId" xmlns:samplens="urn:org.sample.wms"/>'),
#value = #HttpServletRequest.getParameter('oid') == null ?
 'someDefaultValue' : #HttpServletRequest.getParameter('oid') ,
#XmlHelper.setAttribute(#element, 'value', #value)
```

Expected AuthnRequest

A GET request to https://*cpf host>:cpf.https.port>*/sp/startSSO.ping? Partnerldpld=<entityID>&oid=123 would trigger the following Extensions block.

```
<samlp:AuthnRequest ...>
  <saml:Issuer ...>...</saml:Issuer>
  <samlp:Extensions>
    <samplens:orgId name="orgId" value="123"</pre>
xmlns:samplens="urn:org.sample.wms"/>
  </samlp:Extensions>
</samlp:AuthnRequest>
```

Fulfillment by datastore queries

You can configure PingFederate to use local datastores to fulfill various contracts or conditions to verify in the token authorization process.

Datastores represent external systems where user attributes and other data are stored. Once defined, you can configure PingFederate to retrieve user attributes from datastores for contract fulfillment and token authorization in various use cases.

PingFederate supports a wide variety of database servers and directory servers. As needed, the PingFederate SDK supports the creation of custom drivers for connecting to other types of data repositories, such as flat files or SOAP-connected databases. For more information, see the Javadoc for the CustomDataSourceDriver interface, the SamplePropertiesDataStore.java file for a sample implementation, and the SDK Developer's Guide on page 1021 for build and deployment information.

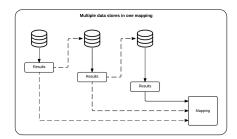


The Javadoc for PingFederate and the sample implementation are in the <pf install>/ pingfederate/sdk directory.

PingFederate can query multiple datastores for additional attributes in most configurations.

Multiple datastores in one mapping

The PingFederate IdP server supports separate datastores to look up attributes for a single mapping. For example, you can query multiple datastores for values and map those values in one mapping, or query a datastore for a value and use that value as input for subsequent queries of other datastores.



If a datastore uses results from previous queries as input, and if the previous queries return no result, PingFederate continues the process by moving on to the next datastore in the setup. If you prefer PingFederate to abort and fail the requests, see Configure the behavior of searching multiple data stores with one mapping.

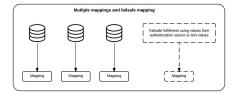
If a required attribute, such as **SAML_SUBJECT** in a SAML contact or **subject** in an authentication policy contract, is not fulfilled, the request fails.

Multiple datastores in one mapping are available for browser single sign-on (SSO) and WS-Trust security token service (STS) on the identity provider (IdP) side, authentication policy contract to service provider (SP) adapter mappings, and the following OAuth workflows:

- Identity provider (IdP) adapter grant mappings
- Resource owner credential grant mappings
- Access token mappings
- OpenID Connect policies (the user-attributes mapping process)

Multiple mappings and failsafe mapping

For browser SSO and WS-Trust STS on the IdP side, PingFederate also supports separate search parameters for each datastore and for "fall-through" searches. For example, you can add the same datastore more than once, using different search queries for each instance, or you can search different datastores successively. If any search fails to find a user in the specified attribute source, the next search is executed until a match is found. You can also configure a failsafe attribute source, providing a default set of attributes from the IdP adapter and using text values.



Datastore query configuration

You can query attributes from a datastore by choosing a datastore instance, specifying the search location and the desired attributes, and entering the search term.

Once the configuration is complete, you can fulfill a contract or verify a condition in the Token Authorization framework using the results from the datastore queries.

This topic describes how to select a datastore for PingFederate to use for looking up attributes.

About this task

On the **Data Store** window, choose a datastore for PingFederate to look up attributes.

Steps

- 1. Enter a description and ID, if prompted, for the datastore.
- 2. Select a datastore instance from the Active Data Store list.



If the datastore you want is not shown in the Active Data Store list, click Manage Data Stores to review or add a datastore instance.

3. Depending on the datastore type, the rest of the setup varies as follows.

Data store type	Required tasks
JDBC	•
	•
LDAP	•
	- (optional)
	•
Other	•

Specifying database tables and columns

Use PingFederate to specify a table and the columns it returns as part of a database query to meet your case needs.

About this task

In the Database Table and Columns window, specify exactly where additional data can be found to fulfill your use case. You can only use one table as a source of data for a database query. For more information about each field, see the following table.

Field	Description
Schema	Lists the table structure that stores information within a database. Some databases require selection of a specific schema for database queries. Other databases do not require selection of a schema.
Table	Displays the tables contained in the database. Select the table to retrieve data from the datastore.
Columns to return from SELECT	Displays selected columns from the selected tables. Select the columns that are associated with the desired attributes you want to return from the database queries.



Important:

For MySQL users — To allow for table and column names that might contain spaces, PingFederate inserts double quotes around the names at runtime. To avoid SQL syntax errors resulting from the quotes, add the session variable sql mode=ANSI QUOTES to the Java Database Connectivity (JDBC) connection string of your datastore instance, as in the following example.

jdbc:mysql://myhost.mydomain.com:3306/pf?sessionVariables=sql mode=ANSI QUOTES

Alternatively, you can configure the system variable sql_mode with the ANSI QUOTES option. For more information, see https://dev.mysql.com/doc/refman/5.7/en/server-system-variables.html.

Steps

- 1. From the **Schema** list, choose a schema when applicable.
- 2. From the **Table** list, select a table.
- 3. Optional: To determine what attributes to examine, click View Attribute Contract.
- 4. Optional: To update an existing configuration where changes to the database might have occurred, click Refresh.
- 5. Under Columns to return from SELECT, choose a column name and click Add Attribute.



You do not have to add a column here to use it as part of a search filter. Add only the column that is required by subsequent sibling configuration items, such as contract fulfillment or token authorization. Any added columns left unused are removed when the configuration is saved.

Repeat this step to add more columns as needed.

Example

Example

Suppose you, the identity provider (IdP), have a data table named **ACCESSTABLE** with thee columns: userid, department, and accesslevel. Your use case requires you to map accesslevel into a SAML contract to an SP.

On the Database Table and Columns window, select the ACCESSTABLE table and add the accesslevel column.

Entering a database search filter

Use a WHERE clause on the **Database Filter** window to set up a database search filter in PingFederate.

About this task

On the Database Filter window, enter a WHERE clause for PingFederate to query the database table you selected to retrieve a record associated with particular values. The clause is in the form:

```
[WHERE] column1=value1
```

The left side (column1) is a column from the database table that you selected on the Database Table and Columns window.



To get a list of columns, click the View List of Columns from ... link.

The right side (value 1) is the match-against value, generally a variable passed in from either an authentication source for an identity provider (IdP) or an assertion for a service provider (SP). The variables are shown underneath the Where text field. If you are retrieving attributes from multiple data You can also apply additional search criteria by using other columns from the targeted table.

Steps

1. Enter a WHERE clause in the text field.



Note:

The initial WHERE is optional.

- 2. Ensure the syntax and variable names are correct. For more information about WHERE clauses, consult your database management system (DBMS) documentation.
- 3. Click **Next** to complete the configuration to query attributes from the database server.

Later in the workflow, you can use the attribute values returned from the database in the applicable contract fulfillment window, the issuance criteria window, or both, to fulfill your use case.

Example

Example

Suppose you have selected a data table named ACCESSTABLE on the Database Table and Columns window. You, the IdP, want to locate user records by matching userid column against the username from an HTML Form Adapter. As a passed-in variable from the HTML Form Adapter, \${username} is shown underneath the Where text field.

On the **Database Filter** window, enter the following filter in the **Where** text field:

```
userid='${username}'
userid
```

The column in the table containing the username information in this example.

'\${username}'

The value of the username variable (username) from an HTML Form Adapter



Important:

You must use the \${{}} syntax to retrieve the value of the enclosed variable and insert single quotation marks around the \${} characters.

Specifying directory properties and attributes

Use these instructions to initiate ways to specify methods for PingFederate to search for particular user data.

About this task

On the LDAP Directory Search window, specify the branch of your directory hierarchy where you want PingFederate to look up user data. For more information about each field, refer to the following table.

Field	Description	
Base DN	The base distinguished name (DN) of the tree structure in which the search begins. This field is optional if records are located at the root of the directory.	

Field	Description
Search Scope	The node depth of the query. Select Subtree (the default value), One level or Object .
Root Object Class	The object class containing the desired attributes.
Attributes	A list of attributes based on the selected Root Object Class value.

Steps

- 1. Optional: Specify a base DN.
- 2. Select a search scope.
- 3. Optional: Click View Attribute Contract to determine what attributes to look up.
- 4. Select a root object class and an attribute, and then click Add Attribute.



You do not have to add an attribute here to use it as part of a search filter. Add only the attributes that are required by subsequent sibling configuration items, such as contract fulfillment or token authorization. Any added attributes that are left unused are removed when the configuration is saved.

Choose from:

Microsoft Active Directory

If you choose the memberof attribute, an optional check box, Nested Groups, appears on the right. Select this check box if you want PingFederate to query for groups the end users belong to directly and indirectly through nested group membership (if any) under the base DN.

For example, if you have three groups under a base DN: Canada, Washington and Seattle. Seattle is a member of Washington. Ana Smith is an end user and a member of Seattle. If the Nested Groups check box is selected, when PingFederate queries for Ana's memberOf attribute values, the expected results are Seattle and Washington. When the Nested Groups check box is not selected (the default), the expected result is Seattle.

Oracle Directory Server or Oracle Unified Directory

Choose isMemberOf under Attribute for nested group membership. For information related to Oracle Directory Server, go to docs.oracle.com/cd/E29127 01/doc.111170/e28967/ ismemberof-5dsat.htm. For information related to Oracle Unified Directory, go to Fusion Middleware Administering Oracle Unified Directory and search for member of user attributes.



If you need to include tokenGroups as one of the attributes, select Object as the search scope and enter a base DN matching the subject DN of the authenticated user—you can use variables from the authentication source (an adapter or an authentication policy contract) or results from the previous lookup in the base DN to fulfill this requirement.

5. Repeat step 4 to add more attributes as needed.

Example

Example

Suppose you want to map the sn Active Directory (AD) user attribute into an OpenID Connect policy. The users for this use case reside under a specific container on your directory server, OU=West, DC=example, DC=com.

On the LDAP Directory Search window, enter OU=West, DC=example, DC=com as the base DN, keep the default Search Scope value (Subtree), select <Show All Attributes> from the Root Object Class list, select the sn AD user attribute, and click Add Attribute.

Defining encoding for binary attributes

Use the LDAP Binary Attribute Encoding Types window to specify an encoding type to apply during fulfillment.

About this task

The LDAP Binary Attribute Encoding Types window appears when at least one attribute is configured as such in the datastore. Because you cannot use binary attribute data in an assertion to the service provider (SP), you must specify the encoding type that you want to apply during fulfillment. The available choices are Base64, Hex, and SID.



Note:

Defining encoding for binary attributes is only applicable to identity provider (IdP) and IdP-to-SP bridging use cases.

Steps

To set an encoding type, select a value from the Attribute Encoding Type list.

Repeat this step for each binary attribute.

Example

Examples

Microsoft Office 365 relies on an immutable Active Directory binary attribute associated with user accounts (objectGUID), and requires this binary data to be Base64-encoded to correlate provisioned federated user data to Active Directory accounts. Select Base64 from the Attribute Encoding Type list.

Claims-based authentication with Microsoft Outlook Web App and Exchange admin center (EAC) requires tokenGroups (another binary attribute in Active Directory) to be SID-encoded. Select SID from the Attribute Encoding Type list.

Entering a directory search filter

Use a filter in PingFederate to query your selected data and retrieve a record associated with it.

About this task

On the LDAP Filter window, enter a filter for PingFederate to query the data you selected. The filter is in the form:

attribute1=value1

The left side (attribute 1) is an attribute from your directory.

Tip:

To see a list of attributes, click the View List of Available LDAP Attributes link.

The right side (value1) is the match-against value, generally a variable passed in from either an authentication source for an identity provider (IdP) or an assertion for a service provider (SP). The variables are shown underneath the Filter text field. If you are retrieving attributes from multiple data stores using one mapping, attributes available from other sources, if previously configured, are listed near the bottom of the window.

You can also apply additional search criteria by using other attributes from the target object class.

A filter narrows a search to locate requested data by either including or excluding specific records. A filter includes the attributes in the search and the value or range of values that the search is attempting to match. Searches are conducted by using three components: at least one attribute (attribute data type) on which to search, a search filter operator that will determine what to match, and the value of the attribute being sought.

Steps

- 1. On the **LDAP Filter** window, enter a search filter in the text field.
- 2. Ensure the syntax and variable names are correct. For general information about search filters, consult your directory documentation.

applicable contract fulfillment window, the issuance criteria window, or both, to fulfill your use case.

3. Click **Next** to complete the configuration to query attributes from the directory server. Later in the workflow, you can use the attribute values returned from your directory server in the

Example

Example

Suppose you want to locate user records by matching the mail Active Directory (AD) user attribute against an extended attribute, eml, in your access token contract for the purpose of mapping attributes to an OpenID Connect policy. As a passed-in variable from the access token, \$ {eml} is shown underneath the Filter text field.

On the LDAP Filter window, enter the following filter in the Filter text field.

```
mail=${eml}
mail
```

An AD user attribute containing the email address of the user

\${eml}

The value of the extended attribute (eml) in the access token contract



Important:

You must use the \${} syntax to retrieve the value of the enclosed variable.

Specifying data source filter and fields

PingFederate allows you to specify a variety of filters, fields, and other settings to best configure your datastore for your organization's needs.

When you choose to retrieve attribute values from other types of datastore, you follow this path through the configuration steps. Windows and steps depend on the implementation of the specific datastore type.

- Specifying a resource path for a REST API datastore on page 369
- Specifying a dynamic authorization header for a REST API datastore on page 369
- Specifying filters and fields for a custom datastore on page 370

Specifying a resource path for a REST API datastore

To set up attribute queries from a REST API datastore, enter the required resource path information.

About this task

PingFederate allows you to specify a resource path for a REST API datastore by entering information in the **Configure Data Source Filters** window.

Steps

On the **Configure Data Source Filters** window, if the REST API datastore requires a relative path or additional query parameters, or both, to retrieve user records, enter them in the **Resource Path** field.

Example

Example

You have use cases that can leverage user attributes obtained through REST APIs. The data source returns user records in JSON. It also provides the following paths to access its data based on user populations:

- https://rest.example.com/development/users
- https://rest.example.com/staging/users

To retrieve the record of a particular user, the request must include the uid query parameter with the identifier of the user.

Your use cases focus on users under the /staging/users path. Your authentication policy uses the HTML Form Adapter, which captures user identifiers using the username attribute.

To address this sample use case:

- 1. Create a REST API datastore with a base URL of https://rest.example.com.
- Add the REST API datastore as an attribute source in the applicable use cases, such as a service provider (SP) connection that uses an HTML Form Adapter instance or an OAuth identity provider (IdP) Adapter Mapping configuration that maps from an HTML Form Adapter instance into the persistent grants.
- 3. When prompted to configure the filtering option on the **Configure Data Source Filters** window, enter /staging/users?uid=\${username} in the **Resource Path** field.

Specifying a dynamic authorization header for a REST API datastore

When you configure an Open ID Connect identity provider (IdP) connection with an application, you can use the access token from the connection as a bearer token in an authorization header to receive additional information as needed.

Before you begin

- Create a Service Provider Open ID Connect IdP connection.
- Configure an Identity Provider authentication policy for the connection.

Steps

- 1. Make the Open ID Connect call to the application to obtain the access token that you plan to use as a bearer token.
 - After you've made the connection, you can find the access token attribute name in <pf install>/ pingfederate/log/server.log in debug mode.
- 2. On the Configure Data Source Filters window, enter the access token attribute name in the Authorization Header field.

Example

Authorization Headers

Authorization Header entries are shown here for Yahoo and Google Open ID Connect IdP connections:

- For Yahoo: Bearer \${idp.https://api.login.yahoo.com.access token}
- For Google: Bearer \${idp.https://accounts.google.com.access token}

Specifying filters and fields for a custom datastore

PingFederate enables you to specify a filter/search query for your data source, then configure fields for your use case.

About this task

Follow this path to set up attribute queries from a custom datastore.

Steps

- 1. On the **Configure Data Source Filters** window, specify a filter, or search query, for your data source.
 - This window display and the syntax of the filter depend on your developer's implementation of the PingFederate SDK.
- 2. On the **Configure Data Source Fields** window, select the fields applicable to your use case.

These choices are supplied by the driver implementation. Select only those needed to fulfill your use case.

Result

You can use the values returned from your data source in the applicable contract fulfillment window, the issuance criteria window, or both.

Configuring failsafe options

From the Failsafe Attribute Source window, you can enable the failsafe mapping or stop the single signon (SSO) transaction when all attribute sources fail to return values for any reason.

About this task

When a datastore is configured and the attribute mappings under Attribute Sources & User Lookup fail to complete the attribute contract in a service provider (SP) connection, you can choose to configure a set of failsafe Attribute Contract Fulfillment mappings. For example, you might configure a set of attributes to identify the SSO subject as a guest user at the SP.



Note:

The Failsafe Attribute Source window does not appear if you have selected the Retrieve additional attributes from multiple data stores using one mapping option on the Mapping Method window.

Important:

The attribute contract is fulfilled using either the mapping configured under Attribute Sources & User Lookup or the failsafe mapping, not both. In other words, you cannot use the failsafe mapping to fill in missing attributes when some are found with the datastore mapping setup but others are not.

The failsafe mapping is used only when all of the mappings configured in the datastore setup fail to return values for any reason. If any mapping succeeds (an attribute mapped to text, for example), failover does not occur.

Alternatively, you can have PingFederate stop the SSO transaction. This choice depends on your agreement with the SP.

Steps

- To enable the failsafe mapping, select Send user to SP using default list of attributes, and then map or enter the desired values on the Attribute Contract Fulfillment window.
- To stop the SSO transaction, select Abort the SSO transaction.

Reviewing datastore query configurations

Review the configuration on the **Summary** window to determine whether to change your configuration, save those changes, or discard them..

Steps

- To amend your configuration, click the corresponding tab title, then follow the configuration wizard to complete the task.
- To keep your changes, click Done and continue with the rest of the configuration.



When editing an existing configuration, you can also click **Save** as soon as the administrative console offers the opportunity to do so.

To discard your changes, click Cancel.

IdP-to-SP bridging

The IdP-to-SP bridging links on the Server Configuration menu provide access to advanced federation settings.

Adapter-to-adapter mappings

PingFederate can act as both identity provider (IdP) and service provider (SP) running on the same server with this configuration.

This configuration is provided for special use cases in which PingFederate is acting as both an IdP and an SP, and user attributes from an IdP adapter are used to create an authenticated session with an SP adapter on the same PingFederate server. Generally, these cases involve software-as-a-service (SaaS) providers who might not support standards-based single sign-on (SSO) but do provide proprietary SSO with "delegated authentication", such as Salesforce and Workday.

In effect, this configuration provides an alternative to setting up complete connections to send SAML assertions and other messages back and forth between an IdP and an SP running on the same PingFederate server in a loop-back configuration to enable nonstandard use cases. Instead, attributes that would normally be sent in an assertion are mapped directly from the IdP authentication adapter to an SP adapter, resulting in a secure SP user session.

To use this configuration, ensure that you have already configured the required IdP and SP adapter instances. You can reuse instances that are also in use for connection configurations.

Managing mappings

Manage adapter-to-adapter mappings using the **Adapter-to-Adapter Mappings** window to control how many you use in PingFederate.

About this task

On the Adapter-to-Adapter Mappings window you can add, modify, or delete adapter-to-adapter mappings. The Adapter-to-Adapter Mappings window is located at Applications # Integration # Adapter-to-Adapter Mappings.

Steps

 To add a mapping, select an identity provider (IdP) adapter instance from the Source Instance list, a service provider (SP) adapter instance from the Target Instance list, and then click Add Mapping.



Note:

You can create only one mapping of a source to the same target. However, you can map different sources to the same target, and the reverse.

- To edit a mapping, select the mapping and then follow the configuration wizard to complete the task.
- To remove a mapping, select **Delete** under **Action** for the mapping.

Assigning a license group

Select a license group if your PingFederate license manages connections by groups.

About this task

Adapter-to-adapter mapping is considered a connection for licensing purposes.



Note:

This window is not displayed for unrestricted or other types of licenses.

Steps

• Select the license group from the list.

Identifying the target application

On the **Target App Info** tab, you can enter the name of the target application and the URL of the application icon.

About this task

The URL is accessible through the IdP Adapter interface, IdpAuthenticationAdapterV2, in the PingFederate Java SDK. For more information about the SDK, see SDK Developer's Guide on page 1021. Both fields are optional.



Note:

If this is a child instance, select the override check box to modify the configuration.

Steps

- 1. Optional: Enter the application name in the **Application Name** field.
- 2. Optional: Enter the URL to the application icon in the Application Icon URL field.
- 3. Click Next.

Configuring attribute sources and user lookup for adapter-to-adapter mappings

To look up attributes and to set up search parameters, configure or modify attribute sources to configure one or more datastores.

About this task

Attribute sources are specific datastore or directory locations containing information possibly required to fulfill the service provider (SP) adapter contract.

Steps

- To configure an attribute source, click Add Attribute Source and complete the setup steps (see Choosing a datastore on page 363).
- To modify an attribute source configuration, select the attribute source and follow the configuration wizard to complete the task.



Note:

Depending on what you change, you might need to modify dependent data in subsequent steps, as indicated.

Configuring target application information

Optionally, you can specify the name and the icon URL of the target application in the Target App Info tab.

Steps

- 1. Enter the target application's name and icon URL.
- 2. Click Next.

Configuring contract fulfillment for adapter-to-adapter mappings

Use one of the listed sources to map each attribute to fulfill the adapter contract.

About this task

The next step in this configuration is to map values from the identity provider (IdP) adapter into the attributes required by the service provider (SP) adapter (the adapter contract).

On the Adapter Contract Fulfillment tab, map each attribute to fulfill the adapter contract from one of these sources:

Adapter

When you make this selection, the associated **Value** drop-down list is populated by the IdP adapter.

Context

Values are returned from the context of the transaction at runtime.



Note:

The HTTP Request selection is retrieved as a Java object rather than text. For this reason, OGNL expressions are more appropriate to evaluate and return values.

Choose Expression and then click Edit to enter an expression (see Using the OGNL edit window on page 220). If the Expression selection is not listed, then the feature is not enabled (see Enabling and disabling expressions on page 215). For syntax and examples, see sections under Construct OGNL expressions on page 216.

LDAP/JDBC/Other (when a datastore is used)

Values are returned from your datastore, if used. When you make this selection, the Value list is populated by the attributes from the datastore.

Expression (when enabled)

This option provides more complex mapping capabilities—for example, transforming incoming values into different formats (see Attribute mapping expressions on page 214). All of the variables available below for text entries are also available for expressions.

No mapping

Select this option to ignore the **Value** field, making the value selection unnecessary.

Text

The value is what you enter. This can be text only, or you can mix text with references to any of the values from the IdP adapter, using the \${attribute} syntax.

The value is what you enter. This can be text only, or you can mix text with references to any of the values from the IdP adapter, using the \${attribute} syntax.

You can also enter values from your datastore, when applicable, using the following syntax.

```
${ds.attr-source-id.attribute}
```

where attr-source-id is the Attribute Source ID value (see Choosing a datastore on page 363) and attribute is any of the datastore attributes you select.

Steps

- 1. Select a source of the attribute under **Source** for each SP adapter contract attribute.
- 2. Specify an attribute under **Value** for each attribute.

Configuring a default target URL (optional)

You can enter a URL in the Default Target URL field to override any SP Default URL SSO setting.

About this task

For more information, see *Configuring default URLs* on page 656.



Note:

If specified, the adapter-to-adapter endpoint parameter tab, enter your default URL into the the TargetResource overrides any default target URL for an adapter-to-adapter mapping. See If specified, the adapter-to-adapter endpoint parameter System-services endpoints on page 1151.

1. On the **Default Target URL**If specified, the adapter-to-adapter endpoint parameter tab, enter your default URL into the the Default Target URL tab.



Note:

This URL overrides any setting within the SP Default URL SSO field.

2. Click Next.

Defining issuance criteria for adapter-to-adapter mappings

You can define the criteria that must be satisfied for PingFederate to further process a request.

About this task

On the **Issuance Criteria** tab, define the criteria to satisfy in order for PingFederate to further process a request. Use this token authorization feature to conditionally approve or reject requests based on individual attributes.

Begin this optional configuration by choosing the source that contains the attribute to verify. Some sources are common to almost all use cases, such as Mapped Attributes. Other sources depend on the type of configuration, such as **JDBC**. Irrelevant sources are automatically hidden. Once you select a source, choose the attribute to verify. Depending on the selected source, the available attributes or properties vary. Specify the comparison condition and the desired value to compare to.

You can define multiple criteria, which must all be satisfied in order for PingFederate to move a request to the next phase. A criterion is satisfied when the runtime value of the selected attribute matches or does not match the specified value, depending on the chosen comparison method. The multi-value contains ... or multi-value does not contain ... comparison methods are intended for attributes that can contain multiple values. Such a criterion is considered satisfied if one of the multiple values match or does not match the specified value. Values are compared verbatim. If you require complex evaluations, including conditional criteria or partial matching, define them using attribute mapping expressions.



Note:

All criteria defined must be satisfied, or evaluated as true, for a request to move forward, regardless of how the criteria were defined. As soon as one criterion fails, PingFederate rejects the request and returns an error message.

Steps

1. On the Issuance Criteria tab, from the Source list, select the attribute's source.

Depending on the selection, the Attribute Name list populates with associated attributes. See the following table for more information.

Source	Description	
Adapter	Select to evaluate attributes from the IdP adapter instance.	

Source	Description	
Context	Select to evaluate properties returned from the context of the transaction at runtime.	
	Note:	
	The HTTP Request context value is retrieved as a Java object rather than text. For this reason, attribute mapping expressions are more appropriate to evaluate and return values.	
JDBC, LDAP, or other types of datastore (if configured)	Select to evaluate attributes returned from a data source.	
Mapped Attributes	Select to evaluate the mapped attributes.	

- 2. From the Attribute Name list, select the attribute to be evaluated.
- 3. From the **Condition** list, select the comparison method.

Available methods:

- equal to
- equal to (case insensitive)
- equal to DN
- not equal to
- not equal to (case insensitive)
- not equal to DN
- multi-value contains
- multi-value contains (case insensitive)
- multi-value contains DN
- multi-value does not contain
- multi-value does not contain (case insensitive)
- multi-value does not contain DN



The first six conditions are intended for single-value attributes. Use one of the multi-value ... conditions for PingFederate to validate whether one of the attribute values matches the specified value. When an attribute has multiple values, using a single-value condition causes the criteria to fail.

4. In the **Value** field, enter the comparison value.



Values are compared verbatim. If you require complex evaluations, including conditional criteria or partial matching, define them using attribute mapping expressions. For more information, see Attribute mapping expressions on page 214.

5. In the **Error Result** field, enter a custom error message.

Error results are handled in one of the two ways.

Redirect

When an InErrorResource URL is provided, the value of the Error Result field is used by the query parameter ErrorDetail in the redirect URL.

Template

When an InErrorResource URL is not provided, the value of the Error Result field is used by the variable \$\mathbb{e}\text{rrorDetail}\$ in the idp.sso.error.page.template.html template file.

Using an error code in the **Error Result** field allows the error template or an application to process the code in a variety of ways, such as displaying an error message or e-mailing an administrator.

To use localized descriptions, enter a unique alias in the Error Result field, such as someIssuanceCriterionFailed. Insert the same alias with the desired localized text in the applicable language resource files, located in the $<pf_install>/pingfederate/server/default/conf/language-packs directory.$

If not defined, PingFederate returns ACCESS DENIED when the criterion fails at runtime.

- 6. Click Add.
- 7. Optional: Repeat to add more criteria.
- 8. If you require complex evaluations, including conditional criteria or partial matching, define them using attribute mapping expressions. For more information, see *Attribute mapping expressions* on page 214.
 - a. Click Show Advanced Criteria.
 - b. In the **Expression** field, enter the required expressions.
 - c. Optional: In the Error Result field, enter an error code or message.



Note:

If the expressions resolve to a string value instead of true or false, the returned value overrides the **Error Result** field value.

- d. Click Add.
- e. Optional: Click **Test**, enter values in the applicable fields, and verify the results.
- f. Optional: Repeat to add multiple criteria using attribute mapping expressions.

Reviewing the adapter-to-adapter mapping

Use the **Summary** tab to edit, save, and discard changes to your configuration of an adapter-to-adapter mapping.

About this task

When you have finished configuring an adapter-to-adapter mapping, you can review the configuration on the **Summary** tab.

Steps

- To keep your changes, click Save.
- To amend your configuration, click the name of the corresponding tab and then follow the configuration wizard to complete the task.
- To discard your changes, click Cancel.

Token translator mappings

Use the token translator configuration to map incoming user attributes from an identity provider (IdP) token processor directly to a service provider (SP) token generator.

This configuration is provided for use cases in which the PingFederate WS-Trust STS exchanges one type of security token for another without needing to generate a SAML token in between. See *About WS-Trust STS* on page 71. Use this configuration, for example, to convert a user's Kerberos token to a third-party proprietary web access management (WAM) session token.

In effect, this configuration provides an alternative to setting up complete STS connections to make such an exchange using the same instance of PingFederate. Instead, incoming user attributes from an IdP token processor are mapped directly to an SP token generator.

To use this configuration, ensure that you have enabled both the IdP and SP roles for PingFederate, including the WS-Trust protocol. See Enabling the WS-Trust protocol on page 951. Make sure to configure the required token-translator instances. You might reuse instances that are also in use for STS connection configurations.

Managing token mappings

Use the **Token Generator Mappings** window to manage your token-to-token mappings.

About this task

On the Applications # Token Exchange # Token Translator Mappings window you can add, modify, or delete token-to-token mappings.

Steps

 To add a mapping, select a token processor instance from the Source Instance list and a token generator instance from the **Target Instance** list, and then click **Add Mapping**.



Note:

You can create only one mapping of a source to the same target. However, you can map different sources to the same target, and vice versa. When you configure multiple identity provider (IdP) token processors, service provider (SP) token generators, or both, you must provide the TokenProcessorId, the TokenGeneratorId query parameters, or both, in the request (see System-services endpoints on page 1151).

- To edit a mapping, select the mapping and then follow the configuration wizard to complete the task.
- To remove a mapping, select **Delete** under **Action** for the mapping.

Configuring attribute sources and user lookup for token mapping

Configure one or more datastores using this optional adapter-to-adapter configuration to look up attributes and to set up search parameters to find information needed to fulfill contracts.

About this task

Attribute sources are specific datastore or directory locations containing information that you might require to fulfill the contract of the token generator.

Steps

- To configure an attribute source, click Add Attribute Source and complete the setup steps. For more information, see *Choosing a datastore* on page 363.
- To modify an attribute source configuration, select the attribute source and follow the configuration wizard to complete the task.



Note:

Depending on what you change, you might need to modify dependent data in subsequent steps, as indicated.

Configuring contract fulfillment for token exchange mapping

Use the listed sources to map values from the token processor into the attributes required by the token generator, the Token Generator Contract.

About this task

Map each attribute to fulfill the Token Generator Contract from one of these sources:

Token

When you make this selection, the associated **Value** drop-down list is populated by the token processor.

LDAP/JDBC/Other (when a datastore is used)

Values are returned from your datastore, if used. When you make this selection, the **Value** list is populated by the attributes from the datastore.

Expression (when enabled)

This option provides more complex mapping capabilities—for example, transforming incoming values into different formats (see *Attribute mapping expressions* on page 214). All of the following variables available for text entries are also available for expressions.

No Mapping

Select this option to ignore the Value field, making it necessary to have no value selection.

Text

The value is what you enter. This can be text only, or you can mix text with references to any of the values from the token processor, using the \${attribute} syntax.

You can also enter values from your datastore, when applicable, using the following syntax.

```
${ds.attr-source-id.attribute}
```

where attr-source-id is the Attribute Source ID value (see Fulfillment by datastore queries on page 361) and attribute is any of the datastore attributes you select.

Steps

- 1. Under **Source**, select a source of the attribute for each attribute in the contract of the token generator.
- 2. Under **Value**, specify an attribute for each attribute.

Defining issuance criteria for token translator mapping

Use the PingFederate issuance criteria features in order to process requests based on individual attributes.

About this task

On the **Issuance Criteria** tab, define the criteria to satisfy in order for PingFederate to further process a request. Use this token authorization feature to conditionally approve or reject requests based on individual attributes.

Begin this optional configuration by choosing the source that contains the attribute to verify. Some sources are common to almost all use cases, such as **Mapped Attributes**. Other sources depend on the type of configuration, such as **JDBC**. Irrelevant sources are automatically hidden. Once you select a source, choose the attribute to verify. Depending on the selected source, the available attributes or properties vary. Specify the comparison condition and the desired value to compare to.

You can define multiple criteria, which must all be satisfied in order for PingFederate to move a request to the next phase. A criterion is satisfied when the runtime value of the selected attribute matches or does not

match the specified value, depending on the chosen comparison method. The multi-value contains ... or multi-value does not contain ... comparison methods are intended for attributes that can contain multiple values. Such a criterion is considered satisfied if one of the multiple values match or does not match the specified value. Values are compared verbatim. If you require complex evaluations, including conditional criteria or partial matching, define them using attribute mapping expressions.



Note:

All criteria defined must be satisfied, or evaluated as true, for a request to move forward, regardless of how the criteria were defined. As soon as one criterion fails, PingFederate rejects the request and returns an error message.

Steps

1. From the **Source** list, select the attribute's source.

Depending on the selection, the Attribute Name list populates with associated attributes. See the following table for more information.

Source	Description
Context	Select to evaluate properties returned from the context of the transaction at runtime.
	Note:
	The HTTP Request and STS SSL Client Certificate Chain context values are retrieved as a Java object rather than text. For this reason, attribute mapping expressions are more appropriate to evaluate and return values.
JDBC, LDAP, or other types of datastore (if configured)	Select to evaluate attributes returned from a data source.
Mapped Attributes	Select to evaluate the mapped attributes.
Token	Select to evaluate attributes from the token processor instance.

2. Select the attribute to evaluate under **Attribute Name**.



To evaluate the STS Basic Authentication Username, STS SSL Client Certificate Chain, or STS SSL Client Certificate's Subject DN context value, ensure that the associated authentication is enabled and configured on System # Server # Protocol Settings to open the WS-Trust STS Settings window.

Available methods:

- equal to
- equal to (case insensitive)
- equal to DN
- not equal to
- not equal to (case insensitive)
- not equal to DN
- multi-value contains
- multi-value contains (case insensitive)
- multi-value contains DN
- multi-value does not contain
- multi-value does not contain (case insensitive)
- multi-value does not contain DN



Note:

The first six conditions are intended for single-value attributes. Use one of the multi-value ... conditions for PingFederate to validate whether one of the attribute values matches the specified value. When an attribute has multiple values, using a single-value condition causes the criteria to fail.



Note:

To evaluate the STS SSL Client Certificate's Subject DN context value, you must select one of the ... DN conditions. These methods normalize the DN before comparison to accommodate for different string representations that are still considered equivalent, such as case sensitivity or white space.

4. In the **Value** field, enter the comparison value.



Note:

Values are compared verbatim. If you require complex evaluations, including conditional criteria or partial matching, define them using attribute mapping expressions. For more information, see Attribute mapping expressions on page 214.

5. In the **Error Result** field, enter a custom error message.

The Error Result field is used by the faultstring element for SOAP 1.1 and the Reason/Text element for SOAP 1.2. For more information on SOAP, see the World Wide Web Consortium's http:// www.w3.org/TR/2000/NOTE-SOAP-20000508/#_Toc478383507.

Using an error code in the Error Result field allows an application to process the code in a variety of ways, such as display an error message or e-mail an administrator.

To use localized descriptions, enter a unique alias in the Error Result field, such as someIssuanceCriterionFailed. Insert the same alias with the desired localized text in the applicable language resource files, located in the <pf install>/pingfederate/server/ default/conf/language-packs directory.

If not defined, PingFederate returns ACCESS DENIED when the criterion fails at runtime.

- 6. Click Add.
- 7. Optional: Repeat to add more criteria.

- 8. If you require complex evaluations, including conditional criteria or partial matching, define them using attribute mapping expressions. For more information, see Attribute mapping expressions on page 214.
 - a. Click Show Advanced Criteria.
 - b. In the **Expression** field, enter the required expressions.
 - c. Optional: In the Error Result field, enter an error code or message.



If the expressions resolve to a string value instead of true or false, the returned value overrides the Error Result field value.

- d. Click Add.
- e. Optional: Click Test, enter values in the applicable fields, and verify the results.
- f. Optional: Repeat to add multiple criteria using attribute mapping expressions.

Reviewing the token exchange mapping

Use the **Summary** window to modify a configuration and then either save, edit, or discard the changes.

About this task

When you finish configuring a token exchange mapping, you can review the configuration on the Summary window.

Steps

- To keep your changes, click Save.
- To amend your configuration, click the name of the corresponding tab and then follow the configuration. wizard to complete the task.
- To discard your changes, click Cancel.

Identity provider SSO configuration

Identity providers (IdP) can use the PingFederate administrative console to configure local applicationintegration information and to manage connections to service provider (SP)-partner sites.

You only need one connection per partner, even if you are targeting more than one web application at the destination SP site.

While your entity ID is defined on the Federation Info tab of the Protocol Settings window, you can identify your organization differently using virtual server IDs on a per-connection basis. For more information, see *Multiple virtual server IDs* on page 106.

You can deploy an SP connection to bridge a service provider to one or more identity providers through authentication policy contracts. For more information, see Federation hub use cases on page 99.



This topic applies to configuration settings needed for browser-based single sign-on (SSO). If you are using PingFederate exclusively as a security token service (STS), start with WS-Trust STS configuration on page 951.

PingFederate uses identity provider (IdP) adapters to look up session information and authenticate users in integrated applications.

The integration of local applications with PingFederate is the essential "first-mile" configuration that allows end-users to access protected resources across domains. This process is facilitated through the use of application-integration kits and a robust SDK (see SSO integration kits and adapters on page 83).

PingFederate uses IdP adapters such as the **HTML Form Adapter** to authenticate users at your site through applications or access-management systems. Go to Authentication # Integration # IdP Adapters to configure instances of IdP adapters.

For authentication sources and selectors that are capable of delegating end-user interactions to external web applications, you can create authentication applications to represent them.

You can also set a default URL to which users can be directed during single logout (SLO), and you can look up system endpoints that application developers at your site need to access PingFederate's single sign-on (SSO) and SLO services.



Note:

If your PingFederate configuration enables the WS-Trust security token service (STS), the selections under Integration also include menu items for configuring plugin Token Processors and optionally STS Request Parameters, see IdP configuration for STS.

Managing IdP adapters

An identity provider (IdP) adapter looks up session information and provides user identification to PingFederate. You must configure at least one instance of an IdP adapter in order to set up connections to service provider (SP) partners.

About this task

PingFederate comes bundled with the PingID integration kit and the following adapters:

- Composite Adapter
- HTML Form Adapter
- HTTP Basic Adapter
- Identifier First Adapter
- Kerberos Adapter
- OpenToken Adapter
- PingID Adapter

Additional integration kits are available on the PingFederate download page under the **Add-ons** tab.

Steps

- 1. Go to Authentication # Integration # IdP Adapters.
- 2. In the Manage IP Adapter Instances window, choose from the following options.

Option	Description
Create New Instance	Configure a new instance
<existing instance="" link=""> under Instance Name</existing>	Modify an existing instance
Check Usage	Review the usage of an existing instance

Option	Description
Delete or Undelete	Remove an existing instance or cancel the removal



Note:

By default, PingFederate automatically checks multi-connection errors whenever you access this window. This verifies that configured connections are not adversely affected by changes made here.

If you experience noticeable delays in accessing this window, you can disable automatic connection validation. Go to System # Server # General Settings.

Creating an IdP adapter instance

The first step in creating an adapter instance is choosing an adapter type.

Steps

- On the **Type** tab, configure the basics of this adapter instance.
 - a. Enter the instance name and ID.
 - b. From the **Type** list, select the adapter type.
 - c. Optional: From the **Parent Instance** list, select an existing type.

If you are creating an instance that is similar to an existing instance, you might consider making it a child instance by specifying a parent. A child instance inherits the configuration of its parent unless overridden. You can specify overrides during the rest of the setup.

Configuring an IdP adapter instance

Use the **IdP Adapters** window to configure existing instances.

About this task

Depending on the selected adapter, the IdP Adapters window presents you with different configuration parameters.

Steps

- 1. Go to Authentication # Integration # IdP Adapters.
- 2. Click the IdP adapter instance you want to configure.
- 3. Follow the in-product instructions to configure the adapter instance.



Note:

If this is a child instance, select the override check box to modify the configuration.

If you are configuring one of the bundled adapters, see Bundled adapters on page 283 or PingID for *PingFederate SSO* for configuration information.

If you are configuring an adapter from an integration kit, including any software as a service (SaaS) connector, go to Integrations Overview, locate the adapter's documentation, and configure the adapter instance accordingly.

Invoking IdP adapter actions

About this task

Adapters can be written to fulfill the following functions:

- Provide configuration assistance
- Perform validation actions
- Generate parameters that might need to be set manually in a configuration file

Steps

Follow the on-screen instructions to complete the actions required.

Extending an IdP adapter contract

You can use the Extended Contract tab to extend the contract of existing IdP adapters, especially those using the Composite Adapter.

About this task

If you are using the Composite Adapter, you must add attributes from the IdP adapter instances that comprise the composite configuration in this tab. For more information, see Composite Adapter on page 283.

If the adapter does not return values for the extended attributes, or if you prefer to fulfill them differently using datastore queries, dynamic text values, or results from OGNL expressions, you can define their fulfillment on the Adapter Contract Mapping tab. For more information, see Defining the IdP adapter contract on page 386.



Note:

If this is a child instance, select the override check box to modify the configuration.

Steps

- 1. Go to Authentication # Integration # IdP Adapters.
- 2. Click the existing IdP adapter instance you want to modify, and then click the Extended Contract tab.
- 3. Enter the name of the desired attribute. Click Add.
- 4. Repeat as needed to add more attributes.
- 5. To save changes, click **Done**.

Setting pseudonym and masking options

Set pseudonym and masking options to uniquely identify a user to your PingFederate SP partners.

Steps

1. Go to Authentication # Integration # IdP Adapters and click the IdP adapter instance that you want to change.



The Override Attributes check box in this window reflects the status of the override option in the Extended Contract tab.

a. Select the check box under Pseudonym for the user identifier of the adapter and optionally for the other attributes, if available.

This selection is used if any of your service provider (SP) partners use pseudonyms for account linking.



Note:

A selection is required whether or not you use pseudonyms for account linking. This allows account linking to be used later without having to delete and reconfigure the adapter. Ensure that you choose at least one attribute that is unique for each user, such as a user's email, to prevent assigning the same pseudonym to multiple users.

- b. Select the check box under Mask Log Values for any attributes that you want PingFederate to mask their values in its logs at runtime.
- c. Select the Mask all OGNL-expression generated log values check box, if OGNL expressions might be used to map derived values into outgoing assertions and you want those values masked.

Defining the IdP adapter contract

You can change the default identity provider (IdP) adapter contract settings using the Adapter Contract Mapping tab.

About this task

An IdP adapter contract is a contract that can be used to fulfill the attribute contract passed to your service provider (SP) partners. By default, PingFederate fulfills the IdP adapter contract with attribute values from the adapter. You can optionally configure PingFederate to fulfill the IdP adapter contract with attribute values from local datastores, dynamic text values, results from OGNL expressions, or a combination of them. In addition, you can verify requests using the Token Authorization framework.

Steps

- 1. Go to Authentication # Integration # IdP Adapters
- 2. Click the **Instance Name** of the existing IdP adapter instance you want to configure.
- 3. Go to the Adapter Contract Mapping tab.



Note:

If this is a child instance, select the **Override Adapter Contract** check box to modify the configuration unless you have already selected the override option in the Extended Contract tab, in which case the Override Adapter Contract check box is automatically selected for you.

- 4. Click Configure Adapter Contract.
 - For information on Attribute Sources & User Lookup, see Defining attribute sources and user lookup on page 387.
 - For information on Adapter Contract Fulfillment, see Configuring IdP adapter contract fulfillment on page 387.
 - For information on **Issuance Criteria**, see *Defining issuance criteria for IdP adapter contract* on page 389.

5. On the Summary tab, click Done to save your adapter contract configurations, or Cancel to discard them.

Defining attribute sources and user lookup

From the Attribute Sources & User Lookup window, you can add new attribute sources or manage existing attribute sources to supply attributes for the identity provider (IdP) adapter contract or the token authorization framework.

About this task

Attribute sources are specific datastore or directory locations containing information that might be needed for the IdP adapter contract or the token authorization framework. You can use more than one attribute source when mapping values to the IdP adapter contract.

For more information about token authorization, see About token authorization on page 96.

The PingFederate IdP server supports separate datastores to look up attributes for a single mapping. For example, you can guery multiple datastores for values and map those values in one mapping, or guery a datastore for a value and use that value as input for subsequent queries of other datastores.

Queries are executed in the order they are displayed on the Attribute Sources & User Lookup tab. Use the up and down arrows as needed to adjust the order.

If a required attribute, such as the user identifier username for the HTML Form Adapter or subject for the OpenToken IdP Adapter, cannot be fulfilled, the request fails.

Steps

- 1. Go to Authentication # Integration # IdP Adapters
- 2. Click the name of the existing IdP adapter instance you want to configure in the **Instance Name** list.
- 3. Click the Adapter Contract Mapping tab.



If this is a child instance, select the Override Adapter Contract check box to modify the configuration unless you have already selected the override option in the Extended Contract tab, in which case the Override Adapter Contract check box is automatically selected for you.

- 4. Click Configure Adapter Contract.
- 5. Choose from:
 - If your use case requires only dynamic texts or results from OGNL expressions without any attributes from local datastores, skip to Configuring IdP adapter contract fulfillment on page 387.
 - To add an attribute source, click Add Attribute Source.
 - To modify an existing instance, select it by name under **Description**.
 - To remove an existing instance or to cancel the removal request, click Delete or Undelete under Action.

Configuring IdP adapter contract fulfillment

You can map values into the IdP adapter contract using the Adapter Contract Fulfillment tab.

Steps

- 1. Go to Authentication # Integration # IdP Adapters
- 2. Click the **Instance Name** of the existing IdP adapter instance you want to configure.



If this is a child instance, select the **Override Adapter Contract** check box to modify the configuration unless you have already selected the override option in the Extended Contract tab, in which case the Override Adapter Contract check box is automatically selected for you.

- 4. Click Configure Adapter Contract.
- 5. Select a source from the **Source** list and specify a **Value** to associate with it.

For more information about the **Source** list and the possible **Values**, see the following table.

Source	Description	
Adapter	Select Adapter to use the attribute value returned by the IdP adapter without customization.	
Context	Select Context to return specific information from the request.	
JDBC, LDAP, or other types of datastores (if	Select an attribute source when PingFederate should retrieve attribute value from a datastore.	
configured)	When you make this selection, the Value list is populated with attributes from your database, directory, or other datastore.	
	Applicable only if you have added at least one attribute source on the Attribute Sources & User Lookup tab. For more information, see <i>Defining attribute sources and user lookup</i> on page 387.	
Expression (if enabled)	Select Expression to support complex mapping requirements; for example, transforming incoming values into different formats. Additionally, the HTTP request is retrieved as a Java object rather than text. For this reason, select Expression as the source and use OGNL expressions to evaluate and return specific information from the HTTP request.	
	Applicable only if you have enabled the use of expressions in PingFederate. For more information, see <i>Attribute mapping expressions</i> on page 214.	
No Mapping	Select No Mapping to ignore the Value field.	
Text	Select Text to return the value you enter under Value .	
	There are a variety of reasons to use a static text value. For example, if the target web application provides a service based on the name of your organization, you may provide the attribute value as a constant.	
	You can mix text with references to attributes from the IdP adapter contract by using the $\{attribute\}$ syntax.	
	You can also enter references to attributes from configured attribute sources by using the \${ds.attr-source-id.attribute} syntax, where attr-source-id is the Attribute Source ID value you entered on Attribute Sources & User Lookup # Data Store and attribute is an attribute from the datastore. For more information, see Defining attribute sources and user lookup on page 387	

- 6. Repeat these steps until all attributes are configured.
- 7. Click Done.

Use the **Issuance Criteria** tab to manage criteria that PingFederate can evaluate to determine whether to issue an identity provider (IdP) adapter contract token for a user.

About this task

On the Issuance Criteria tab, define the criteria to satisfy in order for PingFederate to further process a request. Use this token authorization feature to conditionally approve or reject requests based on individual attributes.

Begin this optional configuration by choosing the source that contains the attribute to verify. Some sources are common to almost all use cases, such as Mapped Attributes. Other sources depend on the type of configuration, such as JDBC. Irrelevant sources are automatically hidden. Once you select a source, choose the attribute to verify. Depending on the selected source, the available attributes or properties vary. Specify the comparison condition and the desired value to compare to.

You can define multiple criteria, which must all be satisfied in order for PingFederate to move a request to the next phase. A criterion is satisfied when the runtime value of the selected attribute matches or does not match the specified value, depending on the chosen comparison method. The multi-value contains ... or multi-value does not contain ... comparison methods are intended for attributes that can contain multiple values. Such a criterion is considered satisfied if one of the multiple values match or does not match the specified value. Values are compared verbatim. If you require complex evaluations, including conditional criteria or partial matching, define them using attribute mapping expressions.



Note:

All criteria defined must be satisfied, or evaluated as true, for a request to move forward, regardless of how the criteria were defined. As soon as one criterion fails, PingFederate rejects the request and returns an error message.

Steps

- 1. Go to Authentication # Integration # IdP Adapters.
- 2. Click the name of the existing instance you want to change from the **Instance Name** list.
- 3. Click Adapter Contract Mapping # Configure Adapter Contract # Issuance Criteria.
- 4. From the **Source** list, select the attribute's source.

Depending on the selection, the Attribute Name list populates with associated attributes. See the following table for more information.

Source	Description	
Adapter	Select to evaluate attributes from the IdP adapter instance.	
Context	Select to evaluate properties returned from the context of the transaction at runtime.	
	Note:	
	As the HTTP Request context value is retrieved as a Java object rather than text, attribute mapping expressions are more appropriate to evaluate and return values.	
JDBC, LDAP, or other types of datastore (if configured)	Select to evaluate attributes returned from a data source.	

- 5. From the Attribute Name list, select the attribute to be evaluated.
- 6. From the **Condition** list, select the comparison method.

Available methods:

- equal to
- equal to (case insensitive)
- equal to DN
- not equal to
- not equal to (case insensitive)
- not equal to DN
- multi-value contains
- multi-value contains (case insensitive)
- multi-value contains DN
- multi-value does not contain
- multi-value does not contain (case insensitive)
- multi-value does not contain DN



Note:

The first six conditions are intended for single-value attributes. Use one of the multi-value ... conditions for PingFederate to validate whether one of the attribute values matches the specified value. When an attribute has multiple values, using a single-value condition causes the criteria to fail.

7. In the Value field, enter the comparison value.



Note:

Values are compared verbatim. If you require complex evaluations, including conditional criteria or partial matching, define them using attribute mapping expressions. For more information, see Attribute mapping expressions on page 214.

8. In the **Error Result** field, enter a custom error message.

To use localized descriptions, enter a unique alias in the Error Result field, such as someIssuanceCriterionFailed. Insert the same alias with the desired localized text in the applicable language resource files, located in the cpf_install>/pingfederate/server/ default/conf/language-packs directory.

If not defined, PingFederate returns ACCESS DENIED when the criterion fails at runtime.

- 9. Click Add.
- 10. Optional: Repeat to add more criteria.

- 11. If you require complex evaluations, including conditional criteria or partial matching, define them using attribute mapping expressions. For more information, see Attribute mapping expressions on page 214.
 - a. Click Show Advanced Criteria.
 - b. In the **Expression** field, enter the required expressions.
 - c. Optional: In the **Error Result** field, enter an error code or message.



Note:

If the expressions resolve to a string value instead of true or false, the returned value overrides the Error Result field value.

- d. Click Add.
- e. Optional: Click **Test**, enter values in the applicable fields, and verify the results.
- f. Optional: Repeat to add multiple criteria using attribute mapping expressions.

Reviewing an IdP adapter contract

From the Summary tab in the Adapter Contract Mapping workflow, review an adapter contract and make changes as needed.

Steps

- To amend your configuration, click the corresponding tab title, then follow the configuration wizard to complete the task.
- To keep your changes, click **Done** and continue with the rest of the configuration.



When editing an existing configuration, you can also click **Save** as soon as the administrative console offers the opportunity to do so.

To discard your changes, click Cancel.

Reviewing and saving an IdP adapter configuration

Review your IdP configuration and save your changes on the **Summary** tab.

Steps

- To keep your changes, click Save.
- To amend your configuration, click the name of the corresponding tab and then follow the configuration. wizard to complete the task.
- To discard your changes, click Cancel.

Authentication applications and the authentication API

The PingFederate authentication API is a JSON-based API that enables end-user interactions, such as credential prompts, to be handled by an external web application. This API does so by providing access to the current state of the flow as an end user steps through a PingFederate authentication policy.

Authentication flows are initiated through browser-based single sign-on (SSO) application endpoints, such as /idp/startSSO.ping, or a protocol request, such as an OpenID Connect authentication request received at the authorization endpoint: /as/authorization.oauth2. As PingFederate runs the configured authentication policy, if it encounters an API-capable adapter or selector, and an authentication application is configured for the policy, PingFederate redirects to the authentication application's URL, passing the ID of the flow in the flowId query parameter.

The authentication application can then retrieve the current state of the flow by issuing a GET request to the /pf-ws/authn/flows/{flowId} endpoint. The links field in the response lists the available

application/vnd.pingidentity.<actionId>+json

When the application invokes an action, PingFederate responds either with the next state for the flow or an

When the user completes the authentication policy steps successfully, the authentication API returns a RESUME status to the authentication application. This status indicates that the API client should redirect the user's browser to the resumeUrl specified in the response. PingFederate will then be responsible for the final step in the flow, such as passing a SAML assertion to a partner. A RESUME status will also be sent if PingFederate encounters an identity provider (IdP) connection in the policy tree, or an IdP adapter or selector that is not API-capable. When the API client redirects the user, PingFederate will take the steps needed to invoke the authentication source.

If the user has interacted with an authentication application and the flow terminates with an error, the API client will receive a FAILED status from the API.

Authentication applications must be highly trusted. See Denying authentication applications access to the authorization endpoint on page 394 for more information on security considerations for authentication applications.



To avoid issues with third-party cookies in some browsers, give the authentication application the same parent domain as the PingFederate base URL.

Key concepts

Flow

The SSO transaction invoking the authentication API.

States

The available states (if any) for a given API-enabled adapter or selector.

Current state

Indicates what the adapter or selector is ready to do next.

Actions

The available actions (if any) for a given state.

Default authentication application

In addition to specifying an authentication application for each policy, you can also configure a default application. PingFederate uses the default application if it encounters an API-capable adapter outside of a policy tree. For example, the default application activates if the user navigates to the Change Password endpoint, /ext/pwdchange/Identify, or the Account Recovery endpoint, /ext/pwdreset/ Identify. PingFederate also uses the default application if authentication policies are disabled, or if the flow falls through to a default authentication source.

You can create and manage authentication applications that use the Authentication API.

About this task

Authentication applications display user interfaces to collect credentials when authentication is completed through the PingFederate authentication API. The default authentication application is used for authentication sources that support the authentication API functionality and are invoked directly, rather than as part of an authentication policy.

Steps

- 1. To manage authentication applications, go to **Authentication # Integration # Authentication API Applications**.
- 2. To toggle the availability of authentication API support, select or clear the **Enable Authentication API** check box.
 - This check box is not selected by default.
- 3. To toggle the availability of the Authentication API Explorer, select or clear the **Enable API Explorer** check box.
 - Applicable and shown if the **Enable Authentication API** check box is selected.
 - When shown, this check box is selected by default.
- 4. Under the **Default Authentication Application** section, perform any of the following actions.

Option	Action
Default Authentication Application	Select an application from the drop-down to designate as the default authentication application.
Check Usage	Click to open a pop-up window listing the configurations in which the authentication is used. Only available for the default authentication application.
Add Authentication Application	Click to add a new authentication. See Configuring authentication applications on page 393.
Delete/Undelete	Click to remove an authentication application or cancel the removal request.

5. To save your changes, click Save.

Configuring authentication applications

Create or configure an authentication application that employs the Authentication API.

Steps

- 1. Go to Authentication # Integration # Authentication API Applications. Choose from:
 - To create a new application, click Add Authentication Application.
 - To modify an existing application, click the **Application Name**.

2. In the **Authentication Application** window, provide information for each field. For more information, see the following table.

Field	Description		
	Description		
Name	The name of the authentication application.		
Description	An optional description of the authentication application.		
URL	The URL of the authentication application.		
Additional Allowed Origins	Enter any number of trusted origins to enable cross-origin resource sharing (CORS) support for the authentication API endpoint.		
	Once configured, client-side web applications from the trusted origins are allowed to make requests to the authentication API endpoint. For more information about CORS, see <i>W3C's recommendation on Cross-Origin Resource Sharing</i> .		
	Sample entry	Expected behavior	
	https:// www.example.c	CORS requests originating from https://	
	https:// www.example.c	CORS requests originating from https://	
	https://www.example.c	CORS requests originating from https://	
		Note:	
		Given this sample entry, a port number is required in the Origin request header.	
	Important:		
	While using the wildcard character provides the convenience of allowing multiple origins with one entry, consider adding individual origins to limit CORS requests to a list of trusted hosts.		
	Click Add to save an entered origin. Repeat to add multiple origins.		

3. To keep your configuration, click **Save** to or click **Cancel** to discard any changes made.

Denying authentication applications access to the authorization endpoint

You can deny authentication applications CORS access to the PingFederate OAuth authorization endpoint.

About this task

Authentication applications must be highly trusted because they have CORS access to the OAuth authorization endpoint /as/authorization.oauth2. They can use an existing session with PingFederate to get tokens for any OAuth client that does not require authentication. Browser-based applications need this level of access to use the *redirectless mode*.

If your deployment does not need this redirectless mode, you can deny authentication applications CORS access to the OAuth authorization endpoint. Applications will still have CORS access to the /pf-ws/authn/flows endpoint but will not be able to directly retrieve OAuth tokens.

- 1. On the administrative console node, open the file authn-api-cors-configuration.xml in the server/default/data/config-store directory.
- 2. Add the following line in the <con:config> section:

```
<con:item name="urlPatterns">/pf-ws/authn/flows(/*)?</con:item>
```

3. Restart PingFederate if it is running as a standalone instance. Otherwise, use the administrative console to replicate the change to the cluster.

Configuring a default URL and error message

As an identity provider (IdP), you can optionally prompt end users to confirm their single logout (SLO) requests and provide a default URL indicating a successful SLO to the end-user, if no other page is designated.

About this task

You can also customize an error message to be displayed as part of the error page rendered in the end-user's browser if an error occurs during IdP-initiated single sign-on (SSO). For example, you might consider modifying the default text to include useful information regarding whom the user should contact or what their next step should be.

Your application or your partner's application can supply the SLO URL at runtime. However, if none is provided, PingFederate will use the default value you enter on this window. For more information, see IdP endpoints on page 1133.

If you leave the default URL blank, PingFederate provides a built-in landing page for the user. This web page is among the templates you can modify with your own branding or other information. For more information, see Customizable user-facing pages on page 798).

Steps

- 1. Go to Authentication # Integration # IdP Default URL.
- 2. Select the check box to prompt the user to confirm SLO.
- 3. Enter a default URL to send the user to on successful SLO.
- 4. Enter a custom error message to display on unsuccessful SLO.



Note:

The error message is displayed only when the application calling the start-SSO endpoint does not explicitly provide its own error page URL. The default entry in this field is used to localize the message. For information about how to use the PingFederate localization feature, see Localizing messages for end users on page 817. If localization is not needed, you can also specify a default message in this field.

5. Click **Save** to save your changes.

Viewing IdP application endpoints

Web-application developers at your site need to know the application endpoints to initiate transactions through PingFederate.

Steps

 Go to System # Endpoints # IdP Application Endpoints to see a list of endpoints and descriptions applicable to your federation role.

These endpoints are built into PingFederate and cannot be changed.

For specific parameters required or allowed for these endpoints, see *IdP endpoints* on page 1133 and *System-services endpoints* on page 1151.

IdP protocol endpoints

PingFederate provides a list of identity provider (IdP) protocol endpoints and exportable metadata for your configuration.

You can find a list of applicable SAML, WS-Federation, and WS-Trust STS endpoints in **System** # **Endpoints** # **IdP Endpoints**. The pop-up window displays only those endpoints related to the federation protocols enabled on **System** # **Server** # **Protocol Settings** # **Federation Info**. These endpoints are built into PingFederate and cannot be changed.

Your federation partners or security token service (STS) clients need to know the applicable IdP services endpoints to communicate with your PingFederate server. Configured service endpoints for SAML connections are included in metadata export files.

PingFederate provides a favicon for all protocol endpoints. For more information, see *Customizing the favicon for application and protocol endpoints* on page 836.

The following table describes each endpoint.

Service	URL and Description	
Single Logout Service (SAML 2.0)	/idp/SLO.saml2	
	The URL that receives and processes logout requests and responses.	
Single Sign-on Service (SAML 2.0)	/idp/SSO.saml2	
	The SAML 2.0 implementation URL that receives authentication requests for processing.	
Artifact Resolution Service (SAML 2.0)	/idp/ARS.ssam12	
	The SOAP endpoint that processes artifacts returned from a federation partner to retrieve the referenced XML message on the back channel. See the note at the end of this table.	
Attribute Query Service (SAML 2.0)	/idp/attrsvc.ssaml2	
	The SAML implementation that receives and processes attribute requests. See the note at the end of this table.	
Single Sign-on Service (SAML 1.x)	/idp/isx.saml1	
	The SAML 1.x implementation of IdP intersite transfer service (ISX) to which clients are redirected for single sign-on (SSO) requests.	



Important:

If mutual SSL/TLS is used for authentication, a secondary PingFederate listening port must be configured and used by partners or STS clients for the relevant endpoints—*.ssaml* and *.wst For more information, see Configuring PingFederate properties on page 752.

Virtual server ID support

For SAML connections using multiple virtual server IDs, each virtual server ID has its own set of protocol endpoints. For more information, see Multiple virtual server IDs on page 106. You can export connection metadata for your partner from **System # Protocol Metadata # Metadata Export**. For more information, see Exporting connection-specific SAML metadata on page 860.

For WS-Federation (and SAML) connections using multiple virtual server IDs, you can provide your partner the federation metadata endpoint, /pf/federation metadata.ping, with the PartnerSpId and vsid parameters, as in the following example.

Partner's entity ID	Your virtual server ID	Federation metadata URL
SP	idev1	https://www.example.com/pf/federation_metadata.ping? PartnerSpId=SP&vsid=idev1

Partner's entity ID	Your virtual server ID	Federation metadata URL
	idev2	https://www.example.com/pf/federation_metadata.ping? PartnerSpId=SP&vsid=idev2

In this example, the base URL and the runtime port of your PingFederate server are www.example.com and 443, respectively.

When the request includes the vsid parameter, the federation metadata endpoint returns information that is specific for a given virtual server ID.

For WS-Trust STS, you can provide your partner the STS metadata endpoint /pf/sts mex.ping with the PartnerSpId and vsid parameters. When the STS metadata request includes the vsid parameter, the STS metadata endpoint returns information that is specific for a given virtual server ID.

For more information about these metadata endpoints, see System-services endpoints on page 1151.



The virtual server ID concept does not apply to the /pf/sts.wst endpoint because token-to-token exchange does not involves any connections. As needed, you can pass the token-to-token endpoint to your partners as-is.

SP connection management

As an identity provider (IdP), you manage connection settings to support the exchange of federationprotocol messages such as SAML, WS-Federation, or WS-Trust with a service provider (SP) or security token service (STS) client application at your site.

The settings your IdP manages include:

- User attributes that you expect to send in a single sign-on (SSO) token, including SAML assertions, WS-Trust STS SAML tokens, or WS-Federation JSON Web Tokens (JWT).
- User attributes that are sent using the SAML Attribute Query profile, if that profile is used. For more information, see Configuring the Attribute Query profile in an SP connection on page 438.
- The protocol, profiles, and bindings of the connection, including detailed security specifications such as the use of back-channel authentication, digital signatures, signature verification, and XML encryption

To establish a connection, you and your partner must decide this information in advance. For more information, see Federation planning checklist on page 104.

If your agreement includes sending assertions containing attribute values from local datastores, you must define the required datastores. For more information, see *Datastores* on page 867.

Administrative interface

Manage connection settings using the SP Affiliations window, accessed from System # Protocol Metadata, which organizes the settings into a series of primary tasks.

Some primary tasks have one or more levels of sub task. Each primary or sub task has its own tab for managing one or more settings. You can move to a sibling task using the Next or Previous button. If you are on a sub task, you can also move to its parent task using the **Done** button.

When creating a new connection, you can save your progress using the Save Draft button. Not all windows offer this option. When you reach the Activation & Summary tab, you must click Save to complete the new connection.

When editing an existing connection, make changes and then click **Save** to commit your changes. You are not required to step through all window to reach the Activation & Summary window before you can save your changes.

Note:

The Save button is available on most tabs. If a window does not show a Save button, click Next or Done until you reach a window where you can use its Save button to commit your changes.

Accessing SP connections

Create, import, and manage service provider (SP) connections.

About this task

The SP Connections window lists all service provider connections. The SP Connections window displays up to 20 connections at a time. As needed, you can use the pagination controls to navigate through your connections, narrow connections by protocol type or status, or search for connection by name or ID.



Tip:

A connection is included in the search results so long as its name or ID is a partial, case-insensitive match to a search term.

Steps

- Go to Applications # Integration # SP Connections.
- To edit a connection, select the connection by its name. For the setting you want to make a change, select the corresponding window title and then follow the configuration wizard to complete the task.
- To create a connection, click Create Connection and follow the on-screen steps.
- To copy a connection, click Select Action # Copy and then follow the on-screen steps.

This is most useful if the new connection and the source connection share many common setting values.



Note:

PingFederate doesn't include outbound provisioning configurations in connection copies. For more information, see Configuring outbound provisioning on page 453.

 To export a connection, click Select Action # Export Connection and then save the XML file as prompted.

This is useful in situations where you want to make a backup of a connection prior to changing it.

 To import a connection, click Import Connection, For more information, see Importing a connection on page 401.

If the connection already exists, you have the option to overwrite the existing connection.



Note:

Prior to the import, you can modify the XML file to suit your needs. The XML file can also be imported to another PingFederate environment acting in the same federation role (IdP) at your site. Note that the source and the target must run the same version of PingFederate.

 To export metadata for any SAML Browser SSO connection, click Select Action # Export Metadata and then follow the on-screen instructions.

 To update a SAML Browser SSO connection, click Select Action # Update with Metadata, then follow the on-screen instructions. For more information, see Importing SP metadata on page 405.

You can update a connection via a metadata XML file or a metadata URL.



Important:

The update operation might require additional configuration. Review the connection after the update operation.

- Click the toggle to enable or disable a connection.
- To remove a connection, click **Select Action** # **Delete**.
- To override the verbosity of runtime transaction logging for all SP connections, click Show Advanced **Fields** and the select the desired override option.

Override option	Description
Off	Select this option and let the per-connection Logging Mode configuration determine the amount of information PingFederate records in the runtime transaction log.
	This is the default selection.
On	Select this option, followed by one of the four logging modes, to set the verbosity of runtime transaction logging for all SP connections. This is most useful when troubleshooting an issue that affects multiple connections.

 To turn off automatic multi-connection error checking, click Show Advanced Fields # Disable Automatic Connection Validation check box.

This check box is not selected by default.

Once selected or cleared, the state of this setting is also reflected on the Authentication # Integration # IdP Connections window.

For more information about this advanced setting and its impact, see Configuring automatic connection validation on page 790.

- To keep your changes, click Save.
- To discard your changes, click Cancel.

Resolving SP connection errors

You can diagnose and resolve service provider (SP) connection errors from the SP Connections menu.

About this task

PingFederate automatically validates configured connections before displaying them on the SP Connections window. This validation ensures these connections have not been adversely affected by any subsequent changes in the supporting components, such as an adapter instance or an authentication policy contract.

If errors are found, the administrative console displays a visual cue next to the applicable connections.

Steps

- 1. Go to Applications # Integration # SP Connections.
- 2. To resolve the error, select the connection and follow the on-screen instructions to modify the configuration, one connection at a time.

Importing a connection

Use the **Import Connection** window to import a connection.

About this task



Prior to the import, you can modify the XML file to suit your needs. The XML file can also be imported to another PingFederate environment acting in the same federation role. The source and the target must run the same version of PingFederate.

Steps

- 1. Go to Applications # Integration # SP Connections.
- 2. Click Import Connection.
- 3. Click **Choose File** and browse to a connection XML file.
- 4. Select the **Allow Update** check box to overwrite an existing connection with the imported file.
- 5. Click **Import** and **Done**.

Updating a SAML connection using metadata

You can update an existing SAML connection using a metadata file or a metadata URL from your partner.

About this task

This manual update is independent from the optional and per-connection automatic update feature.

Steps

- 1. Go to Applications # Integration # SP Connections.
- 2. For the SAML connection that you want to update, click Select Action # Update with Metadata.
- 3. See the following steps to import or update metadata. Instructions vary depending on the medium of the metadata.

Metadata medium Steps	
Metadata XML File	 a. On the Import Metadata window, select the File option. b. Choose the metadata file, and then click Next.
	Note: If the metadata file is digitally signed but the verification certificate is provided outside of the metadata, import the metadata verification certificate on the Import Certificate tab, and then click Next.
	 c. On the Metadata Summary window, review the signature information to evaluate the authenticity of the metadata. d. Click Save.
A metadata URL	a. On the Import Metadata window, select the URL option.

Metadata medium Steps		ps
	b.	Select the metadata from the Metadata URL list.
		Tip:
		If the metadata you want is not shown in the list, click Manage Partner Metadata URLs and enter the URL in the text box.
	c.	Click Load Metadata.
		Note:
		If there is a digital signature error, click Manage Partner Metadata URLs to resolve the issue.
	d.	Click Save.

4. On the Connections window, click Save.

Result



Note:

If the endpoints in the metadata share the same base URL (protocol, hostname, and port), PingFederate uses this information to populate the Base URL field. Consequently, individual endpoints on other windows do not include this information. Only relative paths are shown.

Example

Example

A service provider (SP) has just changed its signing certificate and published a new metadata with the new certificate. To minimize the impacts to your users, you as the identity provider (IdP) can update the SP connection using the metadata immediately.

- 1. Access the SP Connectionswindow from Applications # Integration # SP Connections.
- 2. For the applicable SAML connection, click Select Action # Update with Metadata.
- 3. Follow the workflow to complete the task.

Choosing an SP connection template

The Connection Template tab allows you to choose a quick-connection template for new connection if your installation includes an optional PingFederate SaaS Connector.

About this task



When you select a connection template, many connection settings are configured for you automatically.

- 1. Go to Applications # Integration # SP Connections.
- 2. Click Create Connection.
- 3. To use a template, select Use a template for this connection, then choose the template and enter additional information as required.



Important:

Once you click Next, you cannot return to this window and make a different selection. If you intended to use a different template or no template, you must create a new connection.

Choosing an SP connection type

You can manually create service provider (SP) connections in PingFederate using browser single sign-on (SSO), WS-Trust security token service (STS), outbound provisioning, or any combination thereof.

About this task

If you are not using a connection template, which pre-configures browser-based SSO, indicate on the Connection Type tab whether the connection to this partner is for Browser SSO, WS-Trust STS, outbound provisioning, or any combination of them.



You can add STS, OAuth, and outbound provisioning support to any existing SSO connection, or vice versa, at any time.



Note:

If your partner's deployment supports multiple protocols and you intend to communicate using more than one, you must set up a separate connection for each protocol. Each connection must use a unique (partner) connection ID.

Steps

- 1. Go to Applications # Integration # SP Connections.
- 2. Click Create Connection.
- 3. Select **Do not use a template for this connection**.

4. To configure a connection for secure browser-based SSO, select the Browser SSO Profiles check box.

If you are not using a connection template, you must select the applicable protocol from the list when establishing a new connection.

For a WS-Federation connection, select the desired token type, either SAML 1.1, SAML 2.0, or JWT (JSON Web Token).



Tip:

If you are creating a WS-Federation connection to Microsoft Windows Azure Pack, select JWT as the token type.



PingFederate can encrypt the subject and attributes of SAML 2.0 assertions.

For information about configuring encryption policies on a PingFederate identity provider (IdP), see Configuring XML encryption policy (SAML 2.0).

For information about configuring encryption policies on a PingFederate SP, see Specifying XML encryption policy (for SAML 2.0).

5. Optional: Choose one or both of the following depending on your configuration needs.

Connection Template	Step	
WS-TRUST STS	Select the WS-Trust STS check box.	
Outbound Provisioning	Select Outbound Provisioning and then select the provisioning type from the list.	

6. If your PingFederate license manages connections by groups, select a license group for this connection.

This option is not shown for unrestricted or other types of licenses.

7. To save your settings, click **Next**.

Choosing SP connection options

On the **Connection Options** tab, you can enable browser-based single sign-on (SSO), Attribute Query, or both for the current connection.

Before you begin

For initial steps in creating a service provider (SP) connection, see Choosing an SP connection type on page 403.

Steps

1. Choose one or more of the following options.

Connection option	Description
Browser SSO	Select to create a connection for browser-based SSO.
IdP Discovery	Select to enable identity provider (IdP) discovery. This option is only available if you have configured IdP discovery. For more information,

Connection option	Description
	see Configuring standard IdP Discovery on page 840.
Attribute Query	Select to create a connection that facilitates the SAML 2.0 Attribute Query profile. For more information, see <i>Attribute Query and XASP</i> on page 43.

2. To save your changes, click Next.

Importing SP metadata

When creating or modifying service provider (SP) connections, PingFederate allows you to import metadata from an XML file or a metadata URL.

About this task

If you are using one of the SAML protocols without a connection template, you can expedite the setup by one of the following actions:

- Import a metadata file
- Select a metadata URL

When you select a metadata URL, PingFederate also enables the automatic update option and checks the metadata periodically. If PingFederate detects changes in the partner's signing certificates for digital signature verification, encryption key, or contact information, it updates the connection automatically. For better housekeeping, the update process removes verification certificates from the connection when the partner no longer maintains them in its metadata. In a clustered environment, PingFederate automatically replicates verification certificates and encryption key changes to all engine nodes. Offline engine nodes will also consume these changes as they restart and rejoin the cluster. If you prefer to update the connection manually, you can clear the Enable Automatic Reloading check box.

You can configure reload frequency at System # Protocol Metadata # Metadata Settings # Metadata **Lifetime** tab. The default reload frequency is daily.

We recommend you turn on notifications for SAML metadata update events at System # Monitoring & Notifications # Runtime Notifications.



The notification message provides a list of the applicable items if the metadata contains changes that require additional configuration.

After creating the connection, you can add, remove, or change the metadata URL associated with the connection in the Metadata URL tab. In addition, you can toggle the Enable Automatic Reloading check box for the connection.



Using a metadata URL with automatic reloading streamlines the configuration process. For example, you can quickly establish a browser SSO connection to an InCommon-participating partner. For more information, see www.incommon.org/participants.

1. Select from one of the following steps to import or update metadata.

Metadata medium	Steps
Metadata file	 a. On the Import Metadata tab, select the File option. b. Choose the metadata file, and then click Next.
	Note: If the metadata contains multiple entries, select the desired partner from the Select Entity ID list and click Next.
	Note: If the metadata file is digitally signed but the
	verification certificate is provided outside of the metadata, import the metadata verification certificate on the Import Certificate tab, and then click Next.
	c. On the Metadata Summary tab, review the signature information to evaluate the authenticity of the metadata.
Metadata URL	 a. On the Import Metadata tab, select the URL option. b. Select the metadata from the Metadata URL list.
	Tip: If the metadata you want is not shown in the list, click Manage Partner Metadata URLs. For more information, see Manage Partner metadata URLs on page 629.
	c. Optionally, clear the Enable Automatic Reloading check box to disable automatic update.
	Note: A warning will display if you do not have runtime notifications enabled. To enable these notifications, go to System # Monitoring & Notifications # Runtime Notifications and select the Notification for SAML Metadata Update Events box.

2. Click Next.

Identifying the SP

Add identifying information for service provider (SP) connections that you create or modify.

About this task

On the **General Info** tab, you provide your partner's unique federation identifier, the display name of the connection, and some other optional information, such as virtual server IDs, contact information, and logging mode for runtime transaction logging.

Steps

- 1. For information on initial steps for managing SP connections, see *Choosing an SP connection type* on page 403.
- 2. Provide the basic information to identify your partner.

See the following table for more information.

Field	Description	
Partner's Entity ID,	The published, protocol-dependent, unique identifier of your partner.	
Issuer, Partner's Realm, or Connection	For a SAML 2.0 connection, this is your partner's SAML Entity ID.	
ID (Required)	For a SAML 1.x connection, this is the Audience your partner advertises. This ID may have been obtained out-of-band or using a SAML metadata file.	
	For a WS-Federation connection, this is your partner's Realm.	
	For a security token service (STS)-only connection, you can designate any unique identifier.	
Connection Name	A plain-language identifier for the connection. For example, a company	
(Required)	or department name. This name is displayed in the connection list on the administrative console.	
Virtual Server IDs	If you want to identify your server to this connection partner using an ID other than the one you specified at System # Server # Protocol Settings # Federation Info , enter a virtual server ID in this field and click Add .	
	Enter additional virtual server IDs as needed.	

Field	Description
Base URL	The fully qualified host name and port on which your partner's federation deployment runs. For example, https://www.example.com:9031. This entry is an optional convenience, allowing you to enter relative paths to specific endpoints, instead of full URLs, during the configuration process.
Company	The name of the partner company to which you are connecting.
Contact Name	The contact person at the partner company.
Contact Number	The phone number of the contact person at the partner company.
Contact Email	The email address for the contact person at the partner company.
Application Name	The name of the application, accessible through the IdP Adapter interface IdpAuthenticationAdapterV2 in the PingFederate Java SDK.
	This field is not applicable to an STS-only connection.
Application Icon URL	The URL of the application icon, accessible through the IdP Adapter interface IdpAuthenticationAdapterV2 in the PingFederate Java SDK.
	Note that this field is not applicable to an STS-only connection.
Logging Mode	The level of transaction logging applicable for this connection.

3. After entering the relevant identification information, click **Next**.

Populating extended property values for SP connections

Add, modify, or delete extended properties for service provider (SP) connections on the Extended Properties window.

About this task

Extended property values can serve as metadata. They can also help drive authentication requirements. For more information, see *Extended properties* on page 856.

Steps

- 1. Go to System # Server # Extended Properties.
- 2. Enter the **Name** and **Description** for the property you want to add.
- 3. Optional: Check the **Multivalued** box to indicate that the property permits multiple values.
- 4. Click Add.

Configure IdP Browser SSO

Browser-based single sign-on (SSO), also known as Browser SSO, relies on a user's web browser and HTTP requests to broker identity-federation messaging in XML or JSON web token (JWT) between an identity provider (IdP) and a service provider (SP).

Go to Applications # Integration # SP Connections to access an existing or create a new SP connection. For more information, see Accessing SP connections on page 399.

From the Browser SSO tab inside your SP connection instance, click Configure Browser SSO and follow the steps below based on your federation protocol.



Many steps involved in setting up a federation connection are protocol-independent. They are required steps for all connections, regardless of the associated standards. For more information, see Federation roles on page 28.

Some steps are required under the applicable protocol, while others are optional. Still others are required only based on certain selections. The administrative console determines the required and optional steps based on the protocol and dynamically presents additional requirements or options based on selections.

The following sections provide sequential information about every step you might encounter while configuring browser-based SSO, depending on the protocol you are using for a particular connection.

SAML 2.0 configuration steps

- Choosing SAML 2.0 profiles on page 410
- Setting an SSO token lifetime on page 411
- Configuring SSO token creation on page 411
 - Choosing an identity mapping method for IdP SSO on page 411
 - Setting up an attribute contract on page 414
 - Managing authentication source mappings on page 417
- Configuring protocol settings on page 429
 - Setting Assertion Consumer Service URLs (SAML) on page 430
 - Specifying SLO service URLs (SAML 2.0) on page 434
 - Choosing allowable SAML bindings (SAML 2.0) on page 435
 - Setting an artifact lifetime (SAML) on page 435
 - Specifying artifact resolver locations (SAML 2.0) on page 436
 - Defining signature policy (SAML) on page 436
 - Configuring XML encryption policy (SAML 2.0) on page 437

SAML 1.x configuration steps

- Setting an SSO token lifetime on page 411
- Configuring SSO token creation on page 411
 - Choosing an identity mapping method for IdP SSO on page 411
 - Setting up an attribute contract on page 414
 - Managing authentication source mappings on page 417
- Configuring protocol settings on page 429
 - Setting Assertion Consumer Service URLs (SAML) on page 430
 - Setting a default target URL (SAML 1.x) on page 431
 - Setting an artifact lifetime (SAML) on page 435
 - Defining signature policy (SAML) on page 436

WS-Federation configuration steps

- Setting an SSO token lifetime on page 411
- Configuring SSO token creation on page 411
 - Choosing an identity mapping method for IdP SSO on page 411
 - Setting up an attribute contract on page 414
 - Managing authentication source mappings on page 417
- Configuring protocol settings on page 429
 - Defining a service URL (WS-Federation) on page 432

After configuring SSO settings, you will normally need to configure authentication credentials, the range of which depends on your SSO selection. For more information, see *Configuring credentials* on page 445. You might need to complete further configuration tasks for new or modified connections, depending on the selected options on the **Connection Options** tab.

Choosing SAML 2.0 profiles

On the **SAML Profiles** tab, select one or more SAML 2.0 profiles for your IdP Browser SSO configuration.

About this task

A SAML profile is the message-interchange scenario that you and your federation partner have agreed to use. SAML binding, by contrast, is the transport protocol of SAML messages.



The SAML Profiles tab is not shown for SAML 1.x connections because identity provider (IdP) single sign-on (SSO) is assumed, single logout (SLO) profiles are not supported, and the server supports the "destination-first" (SP-initiated) profile SSO automatically. This window is also not presented for WS-Federation connections because profile selection is not required.



When configuring a local loopback connection, in which one PingFederate instance is both the identity provider and the service provider, disable the IdP-Initiated SLO and SP-Initiated SLO options on the Browser SSO window's SAML Profiles tab. These options determine whether SAML logout requests should be sent to the partner during the SLO flow. Those requests aren't necessary and can cause unexpected behavior when the partner connection exists locally. All local sessions for loopback connections are terminated during the SLO flow without the need to send SAML requests.

For SAML 2.0, PingFederate supports all IdP- and SP-initiated SSO and SLO profiles. For more information on typical SSO and SLO profile configurations, including illustrations, see SAML 2.0 profiles on page 33.

Steps

- 1. Go to Applications # Integration # SP connections.
- 2. Click on the SP connection you want to configure. For more information, see Accessing SP connections on page 399.
- 3. On the Browser SSO tab, click Configure Browser SSO.
- 4. Select either IdP-Initiatied SSO or SP-Initiated SSO or both, depending on your partner agreement. You must select at least one SSO profile.
- 5. Select either IdP-Initiated SLO or SP-Initiated SLO or both, depending on your partner agreement. SLO profile options are only enabled after you choose an SSO profile.

IdP Connections	IdP Connection	Browser SSO

User-Session Creation

A SAML Profile defines what kind of messages may be exchanged between an Identity Provider (IdP) and a Service Provider (SP), and how the messages are transported (bindings). As an SP, you configure this information for your IdP connection.

Summarv

Protocol Settings

Single Sign-On (SSO) Profiles	Single Logout (SLO) Profiles
✓ IDP-INITIATED SSO	✓ IDP-INITIATED SLO
SP-INITIATED SSO	SP-INITIATED SLO

6. Click **Next** to save your changes.

SAML Profiles

Setting an SSO token lifetime

PingFederate gives you the option to change the valid lifetime of the single sign-on (SSO) token.

Before you begin

For previous steps in configuring Browser SSO, see *Choosing SAML 2.0 profiles* on page 410. For more information about managing service provider (SP) connections, see *Accessing SP connections* on page 399.

About this task

Identity-federation standards require a window of time during which a SSO token is considered valid. Each SSO token has an issuance time-stamp element and elements indicating the allowable lifetime of the SSO token before and after the issuance time stamp.

Steps

1. Optional: Override the default values for the following fields.

Field	Description
Minutes Before	The amount of time before the SSO token was issued during which it is to be considered valid.
Minutes After	The amount of time after the SSO token was issued during which it is to be considered valid.

The default value is 5 minutes for both fields.

2. Click Next to save your changes.

Configuring SSO token creation

As an identity provider (IdP), you must specify how PingFederate obtains user-authentication information and use it to create single sign-on (SSO) tokens appropriate for your service provider (SP) partner, including additional user attributes as needed.

About this task

If you are a federation hub bridging a service provider to one or more identity providers, you can associate one or more authentication policy contracts to the SP connection. For more information, see *Federation hub use cases* on page 99.

The configuration involves choosing an identity-mapping method, if applicable; establishing an attribute contract, as needed; and mapping one or more IdP adapter instances, authentication policy contracts, or both.

Steps

- 1. Go to Applications # Integration # SP Connections.
- 2. Click on the SP connection that you want to configure.
- 3. Follow the steps to reach the **Browser SSO** tab for your connection. For more information, see *Configure IdP Browser SSO* on page 408.
- 4. On the Assertion Creation tab, click Configure Assertion Creation.

Choosing an identity mapping method for IdP SSO

In the Identity Mapping window, you choose the type of name identifier your partner requires.

Your selection might affect the way that the service provider (SP) looks up and associates your users at the SP site. You and the SP should decide in advance which option to use.

The choices of name-identifier types depend on whether you use the SAML or WS-Federation protocol. For more information, see one of the following.

- Selecting a SAML Name ID type on page 412
- Selecting a WS-Federation Name ID type on page 413



The Identity Mapping window does not apply for connections using the WS-Federation protocol in conjunction with JSON web token (JWT)-based single sign-on (SSO) tokens. Instead, work with the SP to define an attribute contract that it can use to map users to accounts at the SP site.

Selecting a SAML Name ID type

Choose a name identifier for your SAML Browser single sign-on (SSO) configuration on the identity Mapping tab.

Before you begin

For previous steps in configuring Browser SSO, see Configure IdP Browser SSO on page 408. For more information about managing service provider (SP) connections, see Accessing SP connections on page 399.

About this task

The type of name identifier you select affects how your service provider (SP) partner makes use of account mapping or account linking.

If your SP uses account linking, establishing an attribute contract is not required. However, depending on your agreement, you can choose to supplement the account link with an attribute contract. In this configuration, the account link is used to determine the user's identity, while the additional attributes might be used for authorization decisions, customized web pages, and so on, at the SP site. For more information, see *User attributes* on page 92.



Important:

If you change your configuration to use account linking without additional attributes, any existing attribute contract will be discarded in favor of the new configuration.

Steps

1. Select the type of name identifier that you and your SP have agreed to use.

Option	Description
Standard	Select if you want to send a known attribute to identify a user, for example, a username or an email address.
	In this scenario, the SP often uses account mapping to identify the user locally.

Option	Description
Pseudonym	Select if you and the SP have agreed to use a unique, opaque persistent name identifier, which cannot be traced back to the user's identity at the IdP.
	The SP might also use the identifier for account linking to make a persistent association between the user and a specific local account.
	Select the Include attributes in addition to the pseudonym box if you want to set up an attribute contract to use in conjunction with an opaque identifier. For more information, see <i>Setting up an attribute contract</i> on page 414.
Transient	Select Transient to enhance the privacy of a user's identity. Unlike a pseudonym, a transient identifier is different each time a user initiates SSO.
	An example application for this selection might be when an SP provides generalized group accounts based on organizational rather than individual identity.
	Select the Include attributes in addition to the transient identifier box if you want to set up an attribute contract to use in conjunction with an opaque identifier. For more information, see <i>Setting up an attribute contract</i> on page 414.

2. Click Next to save your changes.

Next steps

If you opted to include attributes in your name identifier, your next step will be to define the attributes. For more information, see Setting up an attribute contract on page 414. Otherwise proceed to Managing authentication source mappings on page 417.

Selecting a WS-Federation Name ID type

Choose a name identifier for your WS-Federation Browser single sign-on (SSO) configuration on the identity Mapping tab.

Before you begin

For previous steps in configuring Browser SSO, see Configure IdP Browser SSO on page 408. For more information about managing service provider (SP) connections, see Accessing SP connections on page 399.

About this task

Your selection might affect the way that the service provider (SP) looks up and associates your users to their local accounts.



Note:

The Identity Mapping window is not applicable to connections using the WS-Federation protocol in conjunction with JSON web token (JWT)-based SSO tokens. Instead, work with the SP to define an attribute contract that it can use to map users to accounts at the SP site.

1. Select the type of name identifier that you and your SP have agreed to use.

Option	Description
Email Address	This attribute is commonly used as a unique identifier for SSO and single logout (SLO). Make this selection, for example, if a user logs in using an email address or if the information is available for lookup in a local datastore.
User Principal Name	The username or other unique ID of the subject initiating the transaction. Make this selection, for example, if a username will be available from the current user session as part of a cookie or can be derived from a local datastore.
Common Name	This selection provides for anonymous SSO to your SP, generally using a hard-coded generalized sign on. Make this selection if your partner agreement involves a many-to-one use case, such as if the SP has a group account set up for all users in a particular domain.

2. Click **Next** to save your changes.

Setting up an attribute contract

Specify the attributes for the name identifier on your WS-Federation or, optionally, for your SAML configuration on the **Attribute Contract** tab.

About this task

An attribute contract is the set of user attributes that you and your partner have agreed will be sent in the single sign-on (SSO) tokens for this connection. For more information, see *Attribute contracts* on page 93.

WS-Federation connections require you to define attribute contracts. For SAML connections, attribute contracts are optional if you are sending either pseudonym or transient identifiers to the partners. For more information, see *Selecting a SAML Name ID type* on page 412.

When establishing an attribute contract, you can change the name format when certain conditions are met. The following table summarizes the conditions and the possible actions that you can perform on the **Attribute Contract** tab.

Protocol	Identity mapping	Attribute contract	SAML_SUBJECT	Additional attributes
SAML 2.0 or SAML	Standard	Required	Built-in.	Optional.
1.1			Subject name format can be changed by selecting a value from a list.	Attribute name format can be changed by selecting a value from a list.
SAML 2.0 or SAML 1.1	Pseudonym or Transient	Required only if the Include	Assumed and cannot be added	At least one is required.
		attributes check box is selected on the Identity Mapping window. Otherwise the Attribute Contract window is not shown.	as an additional attribute.	Attribute name format can be changed by selecting a value from a list.



If you are creating or updating a SAML service provider (SP) connection, consider using the partner's metadata to do so. If the metadata contains the required information, PingFederate automatically populates the attribute contract for you. For more information, see *Importing SP metadata* on page 405.

Steps

- 1. Follow the required steps to create an SSO token depending on your federation protocol. For more information, see Configure IdP Browser SSO on page 408.
- 2. If you are using a SAML protocol, on the Identity Mapping tab you must select either Pseudonym or Transient, and also select the Include Attributes box to access the Attribute Contract tab. For more information, see Selecting a SAML Name ID type on page 412.

3. Optional: Click the Attribute Name Format drop-down to select a different format for the built-in subject identifier, SAML SUBJECT.

Applicable if you and the SP have agreed to a specific format. For more information, see Attribute contracts on page 93.



Note:

As needed, you can customize name-format alternatives in the <pf install>/pingfederate/ server/default/data/config-store/custom-name-formats.xml configuration file. Restart PingFederate to activate any changes made to this file.

- 4. Extend the contract with additional attributes.
 - Enter the name of an additional attribute in the text field under Extend the Contract.

Attribute names are case-sensitive and must correspond to the attribute names expected by your partner.



You can add a special attribute, SAML AUTHN CTX, to indicate to the SP, if required, the type of credentials used to authenticate to the identity provider (IdP) application.

The value of this attribute can then be mapped later on the Attribute Contract Fulfillment window. For more information, see Configuring contract fulfillment for IdP Browser SSO on page 426. The mapped value overrides the authentication context provided by the IdP adapter instance or the Requested AuthN Context Authentication Selector instance, through an authentication policy. If no authentication context is provided by the SAML AUTHN CTX attribute, the IdP adapter instance, or the Requested AuthN Context Authentication Selector instance, PingFederate sets the authentication context as follows:

- For SAML 1.x urn:oasis:names:tc:SAML:1.0:am:unspecified
- For SAML 2.0 urn:oasis:names:tc:SAML:2.0:ac:classes:unspecified



If you are configuring a WS-Federation connection to Microsoft Windows Azure Pack, add upn to the JWT's attribute contract.



If you are configuring a SAML connection to an InCommon participant (see *Incommon* federation participants), the attribute contract might contain or require attributes such as urn:oid:0.9.2342.19200300.100.1.3 and urn:oid:2.5.4.42, which are standard names under various specifications, such as RFC4524 and RFC4519. The following table describes a subset of the object IDs referenced by the most common attributes used by InCommon participants.

Object ID value	Description
0.9.2342.19200300.100.1.3	mail
1.3.6.1.4.1.5923.1.1.1.6	eduPersonPrincipalName
1.3.6.1.4.1.5923.1.1.1.7	eduPersonEntitlement
1.3.6.1.4.1.5923.1.1.1.9	eduPersonScopedAffiliation

Object ID value	Description
1.3.6.1.4.1.5923.1.1.1.10	eduPersonTargetedID
2.5.4.3	cn
2.5.4.4	sn
2.5.4.10	0
2.5.4.42	givenName
2.16.840.1.113730.3.1.241	displayName

For other attributes, see the metadata from your partner. The FriendlyName values, if available, should provide additional information about the attributes. Alternatively, third-party resources, such as https://www.ldap.com/ldap-oid-reference and http://www.oid-info.com/, might help as well.

b. Select an attribute name format from the list.

Applicable if you and the SP have agreed to a specific format. For more information, see Attribute contracts on page 93.



Note:

As needed, you can customize name-format alternatives in the <pf install>/ pingfederate/server/default/data/config-store/custom-name-formats.xml configuration file. Restart PingFederate to activate any changes made to this file.

- c. Click Add.
- d. Repeat until all desired attributes are defined.
- 5. Optional: Click **Edit** to change the configuration of an existing attribute.
- 6. Optional: Click **Delete** to remove an existing attribute.
- 7. Click **Next** to save changes.

Managing authentication source mappings

Use the Authentication Source Mapping tab to map identity provider (IdP) adapters and authentication policies to authenticate users to your service provider (SP).

About this task

IdP adapters are responsible for handling user authentication as part of an single sign-on (SSO) operation. A configured adapter in PingFederate is known as an adapter instance.

In a basic scenario, you map an IdP adapter instance to a SP connection on the Authentication Source Mapping tab and complete its mapping configuration through a series of sub tasks. When a user starts an SSO request, the corresponding IdP adapter is triggered to authenticate the user. Upon successful authentication, PingFederate creates and sends an SSO token to the SP based on the connection settings. As needed, you can map multiple IdP adapter instances to an SP connection, the same IdP adapter instance to multiple SP connections, or a combination of them.

If you use authentication policies to route users through a series of authentication sources and end each successful policy path with an authentication policy contract (APC), you can map the APC to your connection. Like IdP adapter instances, you can map multiple APCs to an SP connection, the same APC to multiple SP connections, or a combination of them.



Zip:

For more information about authentication policies and contracts, see Authentication policies on page 220.

You can also map one or more APCs to an SP connection to bridge a service provider to one or more identity providers. In this scenario, PingFederate is a federation hub for both sides. PingFederate uses APCs to associate this SP connection with the applicable IdP connections to the identity providers. Each APC has its own set of attributes which you map values to the SSO tokens.



Tip:

For more information about the federation hub, see Federation hub use cases on page 99.

Regardless of how many IdP adapter instances and APCs are mapped to an SP connection, PingFederate uses only one adapter instance or policy path to authenticate a user. You can leave the decision to the users or create authentication policies to mandate authentication requirements. Because each adapter instance or APC could return different user attributes, each mapping must define how the attribute contract is fulfilled in its mapping configuration.

Steps

- 1. For initial steps to configure SP connections, see *Accessing SP connections* on page 399.
- 2. For initial steps to configure Browser SSO, see Configure IdP Browser SSO on page 408.
- 3. For initial steps to configure assertion creation, see Configuring SSO token creation on page 411.
- 4. On the **Authentication Source Mapping** tab, select one of the following. Choose from:
 - Click Map New Adapter Instance to map a new IdP Adapter instance. For more information, see Mapping an adapter instance on page 418.
 - Click Map New Authentication Policy to map a new APC. For more information, see Mapping an authentication policy on page 419
 - Click on an existing instance to edit its configuration.
 - Click **Delete** to remove an existing adapter instance or APC. Click **undelete** to cancel the removal request

When authentication sources, such as IdP adapter instances or connection mapping contracts, are restricted to certain virtual server IDs, the allowed IDs are displayed under Virtual Server IDs.

5. When your authentication sources have been mapped, click **Next** save your changes.

Mapping an adapter instance

After extending your attribute contract, the **Authentication Source Mapping** tab gives you the option to map adapter instances.

Steps

- 1. For initial steps, see *Managing authentication source mappings* on page 417.
- 2. On the Authentication Source Mapping tab, click Map New Adapter Instance.
- 3. On the Adapter Instance tab, select an adapter instance from the Adapter Instance drop-down list.

4. Optional: If you want to customize one or more adapter settings for this connection alone, select the Override Instance Settings check box.



Note:

If you are editing a currently mapped adapter instance, you can toggle **Override Instance Settings**. Clearing it removes all previously overridden settings for this connection. Selecting it provides you the opportunity to customize adapter settings specifically for this connection.



Alternatively, you can create child adapter instances of a base adapter instance (with overrides) so that customized settings can be applied to several connections. For more information, see *Hierarchical* plugin configurations on page 90.

Result: Selecting the Override instance settings box will add the Override Instance tab to the navigation bar. For more information, see Overriding an IdP adapter instance on page 419

5. Click **Next**, and refer to the following topics to complete the configuration.

Mapping an authentication policy

After extending your attribute contract, the **Authentication Source Mapping** tab gives you the option to map an authentication policy.

Steps

- 1. Click Map New Authentication Policy.
- 2. On the Authentication Policy Contract tab, select a contract from the Authentication policy contract list.

If you need to create a new contract or manage an existing contract, click Manage Policy Contracts.

3. Click **Next**, and refer to the following topics to complete the configuration.

Overriding an IdP adapter instance

On the **Override Instance** window, you start a series of sub tasks to override adapter settings specifically for this connection.

About this task



Note:

Any changes to the base adapter instance are propagated to a connection as long as you don't override those changes.

Steps

- 1. For initial steps to configure authentication source mapping, see *Managing authentication source* mappings on page 417.
- 2. On the Adapter Instance tab, click Override Instance Settings.

3. On each of the settings tabs, select the **Override** check box, make your changes, and then click **Next**.



If you are editing a currently mapped adapter instance, click Override Instance Settings to reconfigure any overridden settings for this connection. You can also remove all overridden settings on a per-window basis by clearing the **Override** check box near the top of the window.

The override setting windows are functionally identical to those used for creating a new adapter instance. For more information, see *Managing IdP adapters* on page 383.

4. When you are finished, click **Done** to proceed to Selecting an attribute mapping method on page 420.

Restricting an authentication source to certain virtual server IDs

When you multiplex one connection for multiple environments, you can enforce authentication requirements by restricting an authentication source to certain virtual server IDs on the Virtual Server IDs tab.

About this task

Authentication sources are unrestricted by default. For more information, see *Multiple virtual server IDs* on page 106

Steps

- 1. Select the Restrict Virtual Server IDs check box.
- 2. Select one or more virtual server IDs that you want to allow for this authentication source.

Result

If you are editing a currently mapped adapter instance or authentication policy contract (APC), you can toggle the Restrict Virtual Server IDs setting. You can also change the allowed virtual server IDs.

Selecting an attribute mapping method

On the **Mapping Method** tab, you select if and how PingFederate should guery local datastores to help fulfill the attribute contract in conjunction with attribute values from the authentication source.

About this task

To determine whether you need to look up additional values, compare the attribute contract against the adapter contract or the authentication policy contract. If the attribute contract requires more information, you must determine whether local datastores can supply it.



Alternatively, you can configure datastore queries as part of the fulfillment configuration for the applicable identity provider (IdP) adapter contract or authentication policy contract. If so, you do not need to set up datastore query on the connection level.

For more information, see Defining the IdP adapter contract on page 386 or Applying policy contracts or identity profiles to authentication policies on page 250.

Steps

1. For initial steps to configure IdP adapter instances, see *Mapping an adapter instance* on page 418.

2. On the **Mapping Method** tab, select one of the following options.

Mapping method	Description
Retrieve additional attributes from multiple data stores using one mapping	Select to configure one or more datastores to look up attributes for a single mapping.
Retrieve additional attributes from a data store	Select to define alternate datastores to look up attributes and a failsafe mapping configuration. Note: When this option is selected, the token authorization framework, through issuance criteria, does not apply. For more information, see About token authorization on page 96 and Selecting an attribute mapping method.
Use only the adapter contract values in the SAML assertion	Select if you do not require connection-level datastore query.

3. Click **Next** to save changes and proceed to the next tab.

If you opted to require datastore queries, see *Configuring attribute sources and user lookup* on page 425. If not, see *Configuring contract fulfillment for IdP Browser SSO* on page 426.

Configuring default contract fulfillment for IdP Browser SSO

Use the **Attribute Contract Fulfillment** tab to define the default attributes PingFederate will send to the service provider (SP) in case of failure to complete the attribute contract.

Before you begin

For initial steps to configure identity provider (IdP) adapter instances or authentication policy contracts (APC), see *Managing authentication source mappings* on page 417.

If you have selected the failsafe option on the *Mapping Method* tab and the **Send user to SP using default list of attributes** option on the *Failsafe Attribute Source* tab, define the default values that should be sent in the single sign-on (SSO) tokens to the SP.

About this task

On the **Attribute Contract Fulfillment** tab, you must complete the following steps for each adapter instance or APC.

Steps

1. Select a source from the **Source** drop-down list.

- 2. Select a source from the **Source** list and then choose or enter a value. You must map all attributes. See the following table for more information.
 - Adapter or Authentication Policy Contract (the authentication source)

When selected, the Value list is populated with attributes from the authentication source. Select the desired attribute from the list. At runtime, the attribute value from the authentication source is mapped to the value of the attribute in the SSO token.

For example, to map the value of the HTML Form Adapter's username attribute as the value of the SAML SUBJECT attribute on the contract, select Adapter from the Source list and username from the **Value** list.

Context

When selected, the Value list populates with the available context of the transaction. Select the desired context from the list. At runtime, the context value is mapped to the value of the attribute in the SSO token.



Important:

If you are configuring an SP connection to bridge one or more identity providers to a service provider, consider mapping the original issuer of the assertions into an attribute by selecting Context as the source and Authenticating Authority as the value. This is important when bridging multiple identity providers to one service provider, where the service provider should take the information about the original issuer into consideration before granting access to protected resources.

For more information, see *Bridging multiple IdPs to an SP* on page 101.



Note:

Because the HTTP Request context value is retrieved as a Java object rather than text, use OGNL expressions to evaluate and return values (see **Expression**).

Expression

When enabled, this option provides more complex mapping capabilities, such as transforming incoming values into different formats. Select Expression from the Source list, click Edit under Actions, and compose your OGNL expressions. All variables available for text entries are also available for expressions. For more information, see **Text**.

Expressions are not enabled by default. For more information about enabling and editing OGNL expressions, see .

No Mapping

Select this option to ignore the Value field, causing no value selection to be necessary.

Text

When selected, the text you enter is mapped to the value of the attribute in the single sign-on tokens at runtime. You can mix text with references to any of the values from the authentication source using the \${attribute} syntax.



Tip:

Two other text variables are also available: \${SAML SUBJECT} and \${TargetResource}. SAML SUBJECT is the initiating user (or other entity). TargetResource is a reference to the protected application or other resource for which the user requested SSO access; the

\${TargetResource} text variable is available only if specified as a query parameter for the relevant endpoint (either as TargetResource for SAML 2.0 or TARGET for SAML 1.x).

3. After all attributes have been mapped, click **Next** to save changes.

Defining issuance criteria for IdP Browser SSO

Configure the criteria that PingFederate uses to determine user authorization to access service provider (SP) resources.

About this task

On the Issuance Criteria tab, define the criteria that must be satisfied in order for PingFederate to process a request further. This token authorization feature provides the capability to conditionally approve or reject requests based on individual attributes.



Note:

The Issuance Criteria tab does not appear if you have chosen the failsafe option on the Mapping Method tab. For more information, see Selecting an attribute mapping method on page 420.

Begin this optional configuration by adding a criterion. Choose the source that contains the attribute to be verified. Some sources, such as Mapped Attributes, are common to almost all use cases. Other sources, such as JDBC, depend on the type of configuration. Irrelevant sources are automatically hidden. After you select a source, choose the attribute to be verified. Depending on the selected source, the available attributes or properties vary. Finally, specify the comparison method and the desired, compared-to, value.

If you define multiple criteria, all criteria must be satisfied for PingFederate to move a request to the next phase. A criterion is satisfied when the runtime value of the selected attribute matches or does not match the specified value depending on the chosen comparison method. The multi-value contains and multi-value does not contain comparison methods are intended for attributes that might contain multiple values. Such criterion is considered satisfied if one of the multiple values matches or does not match the specified value. Values are compared verbatim. If you require complex evaluations, including conditional criteria or partial matching, define them using attribute mapping expressions.



Important:

When you multiplex one connection for multiple environments, consider using attribute mapping expressions to verify the virtual server ID in conjunction with other conditions, such as group membership information, to protect against unauthorized access. For more information, see Multiple virtual server IDs on page 106 and Issuance criteria and multiple virtual server IDs on page 219.



Note:

All criteria defined must be satisfied or evaluated as true for a request to move forward. As soon as one criterion fails, PingFederate rejects the request and returns an error message.

Steps

1. From the **Source** list, select the attribute's source.

Source	Description
Adapter or Authentication Policy Contract	Select to evaluate attributes from an identity provider (IdP) adapter instance or an authentication policy contract.

Source	Description
Context	Select to evaluate properties returned from the context of the transaction at runtime.
	Note:
	The HTTP Request context value is retrieved as a Java object rather than text. For this reason, attribute mapping expressions are more appropriate to evaluate and return values.
JDBC, LDAP, or other types of datastore (if configured)	Select to evaluate attributes returned from a data source.
Mapped Attributes	Select to evaluate the mapped attributes.

- 2. From the Attribute Name list, select the attribute to be evaluated.
- 3. From the **Condition** list, select the comparison method.

Available methods:

- equal to
- equal to (case insensitive)
- equal to DN
- not equal to
- not equal to (case insensitive)
- not equal to DN
- multi-value contains
- multi-value contains (case insensitive)
- multi-value contains DN
- multi-value does not contain
- multi-value does not contain (case insensitive)
- multi-value does not contain DN



Note:

The first six conditions are intended for single-value attributes. Use one of the multi-value ... conditions for PingFederate to validate whether one of the attribute values matches the specified value. When an attribute has multiple values, using a single-value condition causes the criteria to fail.

4. In the Value field, enter the comparison value.



Note:

Values are compared verbatim. If you require complex evaluations, including conditional criteria or partial matching, define them using attribute mapping expressions. For more information, see Attribute mapping expressions on page 214.

5. In the **Error Result** field, enter a custom error message.

Error results are handled differently for IdP-initiated single sign-on (SSO) and SP-initiated SSO requests.

IdP-initiated SSO

When an InErrorResource URL is provided, the value of the Error Result field is used by an ErrorDetail query parameter in the redirect URL.

Template

When an InErrorResource URL is not provided, the value of the Error Result field is used by the variable \$errorDetail in the idp.sso.error.page.template.html template file.

SP-Intiated SSO

SAML

The Error Result field value is used by the StatusMessage element in the response to the SP.

WS-Federation (Template)

The Error Result field value is used by the \$errorDetail variable in the <pf install>/ pingfederate/server/default/conf/template/sourceid-wsfed-idpexception-template.html template file.

Using an error code in the Error Result field allows the error template or an application to process the code in a variety of ways. For example, the template or application can display an error message or email an administrator.

To use localized descriptions, enter a unique alias in the Error Result field, such as someIssuanceCriterionFailed. Insert the same alias with the desired localized text in the applicable language resource files, located in the <pf install>/pingfederate/server/ default/conf/language-packs directory.

If not defined, PingFederate returns ACCESS DENIED when the criterion fails at runtime.

- 6. Click Add.
- 7. Optional: Repeat to add more criteria.
- 8. If you require complex evaluations, including conditional criteria or partial matching, define them using attribute mapping expressions. For more information, see Attribute mapping expressions on page 214.
 - a. Click Show Advanced Criteria.
 - b. In the **Expression** field, enter the required expressions.
 - c. Optional: In the **Error Result** field, enter an error code or message.



Note:

If the expressions resolve to a string value instead of true or false, the returned value overrides the Error Result field value.

- d. Click Add.
- e. Optional: Click **Test**, enter values in the applicable fields, and verify the results.
- f. Optional: Repeat to add multiple criteria using attribute mapping expressions.

Configuring attribute sources and user lookup

About this task

The order in which attribute sources are listed affects the queries differently based on the selection made on the **Mapping Method** tab. For more information, see *Selecting an attribute mapping method* on page 420.

Retrieve additional attributes from multiple data stores using one mapping

If you plan on using the result of a query as an input to a subsequent query, stack your attribute sources accordingly.

Retrieve additional attributes from a data store

As soon as a query succeeds, PingFederate moves on to the next task, contract fulfillment. Therefore you should prioritize the attribute sources.

Steps

- 1. Click Add Attribute Source and then follow a series of sub tasks to complete the configuration.
- 2. See *Choosing a datastore* on page 363 for instructions on configuring and adding attribute sources.
- 3. Repeat as necessary to add additional sources.

Result

If you are editing a currently mapped adapter instance or authentication policy contract, you can add, remove, or reorder attribute sources, which might require additional configuration changes in subsequent tasks.

Configuring contract fulfillment for IdP Browser SSO

Use the **Attribute Contract Fulfillment** tab to map values to the attributes defined for the contract. These are the values that will be included in the single sign-on (SSO) tokens sent to the service provider (SP).

Before you begin

For initial steps to configure identity provider (IdP) adapter instances or authentication policy contracts (APC), see *Managing authentication source mappings* on page 417.

About this task

If you are bridging one or more identity providers to a service provider, map values to an authentication policy contract. For more information, see *Federation hub use cases* on page 99.

At runtime, an SSO operation fails if PingFederate cannot fulfill the required attribute.

On the **Attribute Contract Fulfillment** tab, you must complete the following steps for each attribute contract.

Steps

1. Select a **Source** from the drop-down.

Adapter or Authentication Policy Contract (the authentication source)

When selected, the Value list is populated with attributes from the authentication source. Select the desired attribute from the list. At runtime, the attribute value from the authentication source is mapped to the value of the attribute in the SSO token.

For example, to map the value of the HTML Form Adapter's username attribute as the value of the SAML SUBJECT attribute on the contract, select Adapter from the Source list and username from the **Value** list.

Context

When selected, the Value list populates with the available context of the transaction. Select the desired context from the list. At runtime, the context value is mapped to the value of the attribute in the SSO token.



Important:

If you are configuring an SP connection to bridge one or more identity providers to a service provider, consider mapping the original issuer of the assertions into an attribute by selecting Context as the source and Authenticating Authority as the value. This is important when bridging multiple identity providers to one service provider, where the service provider should take the information about the original issuer into consideration before granting access to protected resources.

For more information, see *Bridging multiple IdPs to an SP* on page 101.



Note:

Because the HTTP Request context value is retrieved as a Java object rather than text, use OGNL expressions to evaluate and return values (see **Expression**).

LDAP, JDBC, or Other

When selected, the Value list populates with attributes that you have selected in the attribute source configuration. Select the desired attribute from the list. At runtime, the attribute value from the attribute source is mapped to the value of the attribute in the SSO token.

Expression

When enabled, this option provides more complex mapping capabilities, such as transforming incoming values into different formats. Select Expression from the Source list, click Edit under Actions, and compose your OGNL expressions. All variables available for text entries are also available for expressions. For more information, see **Text**.

Expressions are not enabled by default. For more information about enabling and editing OGNL expressions, see .

No Mapping

Select this option to ignore the Value field.

Text

When selected, the text you enter is mapped to the value of the attribute in the single sign-on tokens at runtime. You can mix text with references to any of the values from the authentication source using the \${attribute} syntax.

You can also enter values from your datastore, when applicable, using this syntax:

```
${ds.attr-source-id.attribute}
```

where attr-source-id is the attribute source ID value and attribute is one of the selected attributes in the attribute source configuration.

Note that when using alternate data stores with a failsafe mapping, the attribute source ID value is not applicable. Use the following syntax instead:

\${ds.attribute}



Two other text variables are also available: \${SAML SUBJECT} and \${TargetResource}. SAML SUBJECT is the initiating user (or other entity). TargetResource is a reference to the protected application or other resource for which the user requested SSO access; the \${TargetResource} text variable is available only if specified as a query parameter for the relevant endpoint (either as TargetResource for SAML 2.0 or TARGET for SAML 1.x).

There are also a variety of reasons why you might hard code a text value. For example, if the SP web application provides a service based on the name of your organization, you might provide that attribute value as a constant.

3. After all attributes have been mapped, click **Next** to save changes.



Note:

If you are editing a currently mapped adapter instance or APC, you can update the mapping configuration, which might require additional configuration changes in subsequent tasks.

Reviewing the authentication source mapping

Review and save your authentication source mapping in the Summary tab.

Steps

- To amend your configuration, click the corresponding tab title, then follow the configuration wizard to complete the task.
- To keep your changes, click Done and continue with the rest of the configuration.



When editing an existing configuration, you can also click **Save** as soon as the administrative console offers the opportunity to do so.

To discard your changes, click Cancel.

Reviewing the SSO token creation summary

Review and save your single sign-on (SSO) token creation configuration in the **Summary** tab.

Steps

- To amend your configuration, click the corresponding tab title, then follow the configuration wizard to complete the task.
- To keep your changes, click **Done** and continue with the rest of the configuration.



When editing an existing configuration, you can also click **Save** as soon as the administrative console offers the opportunity to do so.

To discard your changes, click Cancel.

Configuring protocol settings

The Protocol Settings tab provides the launching point for configuring partner endpoints, message customizations, and other protocol-specific settings for browser-based single sign-on (SSO) connections.

Before you begin

For initial steps to configure a service provider (SP) connection, see Accessing SP connections on page

For initial steps to configure Browser SSO, see Configure IdP Browser SSO on page 408.

About this task

SAML 2.0

- Outbound SSO bindings (POST, artifact) and the corresponding assertion consumer service (ACS) URLs
- Outbound SLO bindings (POST, redirect, artifact, SOAP) and the corresponding protocol endpoints
- Inbound bindings (POST, redirect, artifact, SOAP)
- Artifact lifetime
- Signature policy
- Encryption policy

SAML 1.x

- Outbound SSO bindings (POST, artifact) and the corresponding assertion consumer service (ACS) URLs
- Default target URL
- Artifact lifetime
- Signature policy

WS-Federation

- Protocol endpoint
- Default target URL

Steps

1. Before configuring Browser SSO protocol settings, you must first configure assertion configuration. For more information, see Configuring SSO token creation on page 411

2. In the Protocol Settings tab, click Configure Protocol Settings to begin.

Setting Assertion Consumer Service URLs (SAML)

If your PingFederate configuration uses any version of SAML, you can configure assertion indexes, bindings, and endpoint URLs on the Assertion Consumer Service URL tab.

Before you begin

For prerequisites and initial steps to configure Browser SSO protocol settings, see Configuring protocol settings on page 429.

About this task

The assertion consumer service (ACS) endpoint is a location to which the single sign-on (SSO) tokens are sent, according to partner requirements. ACS is applicable to all SAML versions and both the identity provider (IdP)- and service provider (SP)-initiated SSO profiles.



The SP might request that the SAML assertion be sent to one of several URLs, using different bindings. PingFederate uses the defined URL entries on this page to validate the authentication request. However, per SAML specifications, if the request is signed, PingFederate can verify the signature instead. The ACS URL does not necessarily need to be listed here. This is useful for scenarios where an ACS URL might be dynamically generated.

Some federation use cases might require additional customizations in the assertions sent from the PingFederate IdP server to the SP, such as placing well-formed XML in the <attributeValue> element or including the optional SessionNotOnOrAfter attribute in the <AuthnStatement> element. You can use OGNL expressions to fulfill these use cases.

Steps

- 1. In the Assertion Consumer Service URL tab, configure one or more SAML ACS endpoints.
 - a. Select a SAML binding from the **Binding**drop-down list.
 - b. Enter the ACS endpoint URL to the **Endpoint URL** field.

You can enter a relative path (begin with a forward slash) if you have provided a base URL on the General Info window.

Optional: Select the **Default** box if you want this entry to be the default ACS endpoint.

The administrative console always sets the first entry as the default ACS endpoint. You can reset the default endpoint when you add ACS endpoint.

d. Optional: Enter an integer to the **Index** field for this ACS endpoint.

The administrative console automatically assigns an index value for each ACS endpoint, starting from 0. If you want to define your own index values, you must make sure the index values are unique.

- e. Click Add.
- f. Optional: Repeat to add additional ACS endpoints.

2. Optional: Customize messages using OGNL expressions.



OGNL expressions are not enabled by default. For more information about enabling and editing OGNL expressions, see Attribute mapping expressions on page 214.

- Click Show Advanced Customizations.
- Select a message type from the list.
- c. Enter an OGNL expression to fulfill your use case.



For more information about **Message Type**, available variables, and sample OGNL expressions, see Customizing assertions and authentication requests on page 355.

- d. Click Add.
- e. Optional: Repeat to add another message customization.
- 3. Click Next to proceed to the next tab. For SAML 1.x configurations, see Setting a default target URL (SAML 1.x) on page 431. For SAML 2.0, see Specifying SLO service URLs (SAML 2.0) on page 434.

Result

If you are editing an existing connection, you can reconfigure any items, which could require additional configuration changes in subsequent tasks. You must always configure at least one ACS endpoint.

Setting a default target URL (SAML 1.x)

SAML 1.x service provider (SP) connections requires that a default target URL be specified for a scenario where the identity provide (IdP) application does not include one in its single sign-on (SSO) request. This default URL represents the destination on the SP where the user will be directed.

Before you begin

For prerequisites and initial steps for configuring Browser SSO protocols, see Configuring protocol settings on page 429.

Steps

- 1. Enter default destination in the **Default Target URL** field.
- Click Next to save your setting.

Result

If you are editing an existing connection, you update the target destination. You must always define a default destination when configuring a SAML 1.x SP connection.

Specifying the WS-Trust version

You can specify whether to use WS-Trust version 1.2 or 1.3 for tokens. The default version is 1.2.

Before you begin

For prerequisites and initial steps for configuring Browser SSO protocols, see Configuring protocol settings on page 429.

1. From the **WS-Trust Version** drop-down list, select version **1.2** or **1.3**.



For version 1.3, the response is always a RequestSecurityTokenResponseCollection object, as in the following example.

<wst:RequestSecurityTokenResonseCollection xmlns:wst="http://docs.oasis-</pre> open.org/ws-sx/ws-trust/200512/">

For more information about WS-Trust version 1.3, see OASIS WS-Trust 1.3 Standards.

2. Click **Next** to save your changes.

Defining a service URL (WS-Federation)

On the Service URL tab, enter the WS-Federation protocol endpoint of your service provider (SP) partner where PingFederate will send single sign-on (SSO) tokens and single logout (SLO) cleanup messages.

Before you begin

For prerequisites and initial steps for configuring Browser SSO protocols, see Configuring protocol settings on page 429.

About this task

The SSO tokens are transmitted within a Request for Security Token Response (RSTR) message in response to a request for authentication from the SP. SLO cleanup messages are sent to your partner when PingFederate, as the identity provider (IdP), receives a user's SLO request. These cleanup messages indicate that the user's local session has been terminated.

To protect against session token hijacking, you can specify additional allowed domains and paths on this window. If the option to validate wreply for SLO is enabled, these additional domains and paths will also be taken into consideration as well. For more information, see Managing partner redirect validation on page 634.

Some federation use cases might require additional customizations in the RSTR message sent from the PingFederate IdP server to the SP. You can use OGNL expressions to fulfill these use cases.

Steps

1. On the Service URL tab, enter the WS-Federation protocol endpoint at the SP site in the Endpoint **URL** field.

You can enter a relative path (begin with a forward slash) if you have provided a base URL on the **General Info** tab. For more information, see *Identifying the SP* on page 407.

- 2. Optional: Specify additional allowed domains and paths.
 - Indicate whether to mandate secure connections when this resource is requested under Require HTTPS.



Important:

This selection is recommended to ensure that the validation will always prevent message interception for this type of potential attack, under all conceivable permutations.

This check box is selected by default.

b. Enter the expected domain name or IP address of this resource under Valid Domain Name.

Enter a value without the protocol, such as example.com or 10.10.10.10.

Prefix a domain name with a wildcard followed by a period to include subdomains using one entry. For instance, *.example.com covers hr.example.com or email.example.com but not example.com, the parent domain.



Important:

While using an initial wildcard provides the convenience of allowing multiple subdomains using one entry, consider adding individual subdomains to limit the redirection to a list of known hosts.

c. Optional: Enter the exact path of this resource under Valid Path.

Start with a forward slash, without any wildcard characters in the path. If left blank, any path under the specified domain or IP address is allowed. This value is case-sensitive. For instance, /inbound/Consumer.jsp allows /inbound/Consumer.jsp but rejects /inbound/consumer.jsp.

You can allow specific query parameters with or without a fragment by appending them to the path. For instance, /inbound/Consumer.jsp?area=West&team=IT#ref1001 matches /inbound/Consumer.jsp?area=West&team=IT#ref1001 but not /inbound/Consumer.jsp?area=East&team=IT#ref1001.

d. Optional: Select the check box under **Allow Any Query/Fragment** to allow any query parameters or fragment for this resource.

Selecting this check box also means that no query parameter and fragment are allowed in the path defined under **Valid Path**.

This check box is not selected by default.

e. Click Add.

Use the **Edit**, **Update**, and **Cancel** workflow to make or undo a change to an existing entry. Use the **Delete** and **Undelete** workflow to remove an existing entry or cancel the removal request.

f. Repeat these steps to define multiple expected resources.



Note:

The display order does not matter. A more specific match is considered a better match and an exact match is considered the best match.

3. Optional: Customize messages using OGNL expressions.

Expressions are not enabled by default. For more information about enabling and editing OGNL expressions, see Attribute mapping expressions on page 214.

- a. Click Show Advanced Customizations.
- b. Select a message type from the list.
- c. Enter an OGNL expression to fulfill your use case.



For more information about **Message Type**, available variables, and sample OGNL expressions, see Customizing assertions and authentication requests on page 355.

- d. Click Add.
- e. Optional: Repeat to add another message customization.
- 4. Click **Next** to save your changes.

Result

If you are editing an existing connection, you can reconfigure any items, which might require additional configuration changes in subsequent tasks.

Specifying SLO service URLs (SAML 2.0)

On the SLO Service URLs tab, you associate bindings to the endpoints where your service provider (SP) receives logout requests when single logout (SLO) is initiated at your site and where PingFederate sends SLO responses when it receives SLO requests from the SP.

Before you begin

For prerequisites and initial steps for configuring Browser SSO protocols, see Configuring protocol settings on page 429.

About this task

This step applies only to SAML 2.0 connections when either SLO profile is selected on the SAML Profiles tab. For more information, see *Choosing SAML 2.0 profiles* on page 410.

Steps

- 1. Select a SAML binding from the list.
- 2. Enter the SLO endpoint URL to the Endpoint URL field.

You can enter a relative path (begin with a forward slash) if you have provided a base URL on the **General Info** tab. For more information, see *Identifying the SP* on page 407.

3. Optional: Enter a URL in the Response URL field.

When specified, this URL is the location to which SLO logout response messages are sent based on your partner agreement. When omitted, PingFederate sends logout responses to the SLO endpoint URL.

You can enter a relative path (begin with a forward slash) if you have provided a base URL on the **General Info** window. For more information, see *Identifying the SP* on page 407.

- 4. Click Add.
- 5. Optional: Repeat to add additional SLO endpoints.
- 6. Click **Next** to save your settings.

If you are editing an existing connection, you can reconfigure the SLO endpoints, which might require additional configuration changes in subsequent tasks.

Choosing allowable SAML bindings (SAML 2.0)

On the **Allowable SAML Bindings** tab, you select the one or more bindings that your service provider (SP) partner can use to send SAML authentication requests or single logout (SLO) messages.

Before you begin

For prerequisites and initial steps for configuring Browser SSO protocols, see *Configuring protocol settings* on page 429.

About this task

This step applies only to SAML 2.0 connections when the SP-initiated SSO profile or either SLO profile is selected on the **SAML Profiles** tab.

Steps

1. On the **Allowable SAML Bindings** tab, select the applicable SAML bindings based on your partner agreement.



Note:

If you have specified an Assertion Consumer Service (ACS) or SLO endpoint using the artifact (outbound) binding, you must including SOAP as one of the allowable (inbound) binding.

2. Click **Next** to save changes and proceed to **Artifact Resolver Locations**. For more information, see *Specifying artifact resolver locations (SAML 2.0)* on page 436.

Result

If you are editing an existing connection, you can reconfigure the allowable bindings, which might require additional configuration changes in subsequent tasks.

Setting an artifact lifetime (SAML)

When PingFederate sends an artifact to your service provider (SP)'s SAML ACS endpoint or SAML 2.0 SLO endpoint, an element in the message indicates how long it should be considered valid. On the **Artifact Lifetime** tab, specify the expiry information in seconds.

Before you begin

For prerequisites and initial steps for configuring Browser SSO protocols, see *Configuring protocol settings* on page 429.

About this task

You can change the default value to meet your requirements. You should also consider synchronizing your serve clock with your partner's SAML gateway server. If clocks are not synchronized, you might need to set the artifact lifetime to a higher value to prevent latency issues.

Steps

- Optional: On the Artifact Lifetime tab, override the default value of the Artifact Lifetime field.
 The default value is 60 (seconds).
- 2. Click Next to save your changes.

Specifying artifact resolver locations (SAML 2.0)

When the artifact binding is enabled as one of the allowable bindings on the Allowable SAML Bindings tab, you must provide at least one artifact resolution service (ARS) endpoint on the Artifact Resolver Locations tab.

About this task

The ARD endpoint is where PingFederate sends back-channel requests to resolve artifacts received from the service provider (SP).

Steps

- 1. On the Artifact Resolver Locations tab, enter the ARS endpoint URL in the URL field.
 - You can enter a relative path (begin with a forward slash) if you have provided a base URL on the **General Info** tab. For more information, see *Identifying the SP* on page 407.
- 2. Optional: Enter an integer to the **Index** field for this ACS endpoint.
 - The administrative console automatically assigns an index value for each ARS endpoint, beginning with 0. If you want to define your own index values, you must make sure the index values are unique.
- Click Add.
- 4. Optional: Repeat to add additional ARS endpoints.



Note:

When specifying multiple ARS endpoints, each endpoint must share the same transport protocol. That is, if one endpoint uses HTTPS, then all must use HTTPS.

5. After you have entered all of your ARS endpoints, click **Next** to save changes.

Result

If you are editing an existing connection, you can reconfigure any ARS endpoints.

Defining signature policy (SAML)

Use the Signature Policy tab to control how digital signatures are used for SAML messages.

Before you begin

For prerequisites and initial steps for configuring Browser SSO protocols, see Configuring protocol settings on page 429.

About this task

The choices made in this tab depend on your partner agreement and your federation protocol. For more information, see Digital signing policy coordination on page 88.

SAML 2.0

Digital signing is required for SAML response messages sent from the identity provider (IdP) with the POST or redirect binding. Based on the SAML specifications, PingFederate provides three options:

- Select Always Sign Assertion to always sign the assertion portion inside the SAML response
- Select Sign Response As Required to sign the SAML response message per the SAML specifications. This is the default selection.
- Select both to always sign the assertion portion inside the SAML response message for all bindings and to sign the SAML response message per the SAML specifications.

Authentication request messages from the service provider (SP) may also be signed to enforce security. This scenario applies only when the SP-initiated single sign-on (SSO) profile is enabled on the SAML Profiles tab. SelectRequire Authn Requests to be Signed to enforce this digital signature requirement. For more information, see Choosing SAML 2.0 profiles on page 410.

SAML 1.x

For SAML 1.0 and SAML 1.1, the assertion portion inside the SAML response message can be digitally signed.

 Select Always Sign Assertion to always sign the assertion portion inside the SAML response message.

Steps

- 1. On the **Signature Policy** tab, select the options based on your partner agreement and federation protocol.
- 2. Click Next to save changes.

Result

If you are editing an existing connection, you can reconfigure the digital signature policy, which might require additional configuration changes in subsequent tasks.

Configuring XML encryption policy (SAML 2.0)

For SAML 2.0 configurations, in addition to using signed assertions to ensure authenticity, you and your partner can also agree to encrypt all or part of an assertion to improve privacy. If so, you can configure these settings on the Encryption Policy tab.

Before you begin

For prerequisites and initial steps for configuring Browser SSO protocols, see Configuring protocol settings on page 429.

About this task



Note:

For WS-Fed connections with SAML 2.0 assertions, you cannot encrypt the entire assertion.

Option	Name identifier (SAML_SUBJEC		Encrypt the SAML_SUBJECT in SLO messages to the SP	Allow encryption in SLO messages from the SP
None	No encryption.	No encryption.	No encryption.	No encryption.
The entire assertion	Encrypted.	Encrypted.	Available as an option.	Available as an option.
One or more attributes	Available as an option.	Available as an option.	Available as an option only if you select to encrypt the name identifier (SAML_SUBJECT).	Available as an option only if you select to encrypt the name identifier (SAML_SUBJECT).

Steps

1. Select the options based on your partner agreement.

Result

If you are editing an existing connection, you can reconfigure the XML encryption policy, which might require additional configuration changes in subsequent tasks.

Reviewing protocol settings

Review and save your protocol settings on the **Summary** tab.

Steps

- To amend your configuration, click the corresponding tab title, then follow the configuration wizard to complete the task.
- To keep your changes, click **Done** and continue with the rest of the configuration.



When editing an existing configuration, you can also click Save as soon as the administrative console offers the opportunity to do so.

To discard your changes, click Cancel.

Reviewing browser-based SSO settings

Review and save your browser-based single sign-on (SSO) settings on the **Summary** tab.

Steps

- To amend your configuration, click the corresponding tab title, then follow the configuration wizard to complete the task.
- To keep your changes, click Done and continue with the rest of the configuration.



When editing an existing configuration, you can also click Save as soon as the administrative console offers the opportunity to do so.

To discard your changes, click Cancel.

Configuring the Attribute Query profile in an SP connection

At the Attribute Query step, you configure your connection to respond to requests for user attributes from your partner service provider (SP), if you have chosen this option.

Before you begin

For prerequisites and previous steps to configure identity provider (IdP) Browser single sign-on (SSO), see Configure IdP Browser SSO on page 408.

About this task

For more information about the Attribute Query profile, see Choosing SP connection options on page 404.

Attribute queries do not depend on SSO, but can be used independently or in conjunction with Browser SSO or provisioning to provide flexibility in how a user authenticates with SP applications. For more information, see Attribute Query and XASP on page 43.

Steps

On the Attribute Query tab, click Configure Attribute Query Profile. See Defining retrievable attributes on page 439 for next steps.

Defining retrievable attributes

On the Retrievable Attributes window, you specify the user attributes you and your partner have agreed to allow in an attribute query transaction.

Before you begin

For prerequisites and previous steps to configure the Attribute Query profile, see Configuring the Attribute Query profile in an SP connection on page 438.

About this task



The service provider (SP) might not necessarily request all of these attributes in each attribute-query request. Instead, the list simply limits the request to a subset of these attributes.

Steps

1. On the **Retrievable Attributes** tab, follow these steps to configure your Attribute Query attributes. To add an attribute, enter the attribute name in the text box and then click Add.



Attribute names are case-sensitive and must correspond to the attribute names expected by your partner.

2. To modify an attribute name:

Action	Steps
Add an attribute	Enter the attribute name in the text box. Click Add .
Modify an existing attribute	Click Edit . Modify the attribute name in the text box. Click Update .
Delete an existing attribute	Click Delete.

3. Click **Next** to save your changes.

Configuring attribute lookup

The optional attribute lookup configuration allows you to configure one or more datastores to look up attributes and to set up search parameters.

Before you begin

For prerequisites and previous steps to configure the Attribute Query profile, see Configuring the Attribute Query profile in an SP connection on page 438.

About this task

Attribute sources are specific datastore or directory locations containing information that is returned to the service provider (SP) in response to an attribute request.

- 1. On the Attribute Sources & User Lookup tab, you can do the following Choose from:
 - To configure an attribute source, click Add Attribute Source and complete the setup steps. For more information, see *Choosing a datastore for Attribute Query* on page 440.
 - To modify an attribute source configuration, select the attribute source and complete the setup steps.



Note:

Depending on what you change, you might need to modify dependent data in subsequent steps, as indicated.

2. When your attribute sources are configured, click **Next** to save your changes.

Choosing a datastore for Attribute Query

On the **Data Store** tab, choose a datastore instance for PingFederate to look up attributes.

Before you begin

For prerequisites and previous steps to configure the Attribute Query profile, see Configuring the Attribute Query profile in an SP connection on page 438.

About this task

The process of configuring PingFederate to look up attributes in a datastore for attribute-query responses is similar to that used for single sign-on (SSO) Attribute Sources and User Lookup.

Steps

- 1. Enter a **Description** for the datastore in the text box.
 - a. If prompted, enter an ID in the text box.
- 2. Select a datastore instance from the **Active Data Store** list.



If the datastore you want is not shown in the Active Data Store list, click Manage Data Stores to review or add a datastore instance. For more information, see *Datastores* on page 867.

3. Depending on the datastore type, the rest of the setup varies as follows.

Data store type	Required tasks
JDBC	 Specifying database tables and columns on page 363 Entering a database search filter on page 364
LDAP	 Specifying directory properties and attributes on page 365 Defining encoding for binary attributes on page 367 (optional) Entering a directory search filter on page 367

Data store type	Required tasks
Other	Specifying data source filter and fields on page 369



Important:

When attribute queries are sent using X.509 Attribute Sharing Profile (XASP), use the variable \${SubjectDN}—rather than \${SAML SUBJECT}—to retrieve the subject identifier.

You can also use any of these distinguished name (DN)-parsing variables:

- \${CN}
- \${OU}
- \${0}
- \${L}
- \${S}
- \${C}
- \${DC}

If more than one value exists for any of the parsing variables, then they are enumerated. For example, if the Subject DN is cn=John Smith, ou=service, ou=employee, then you could use any of these elements in your filter qualifier:

- \${SubjectDN}=cn=John Smith,ou=service,ou=employee
- \${ou}=service
- \${ou1}=employee

For more information about XASP, see Attribute Query and XASP.

4. When you have finished configuring your datastore, click **Next** to save changes.

Configuring mapping fulfillment for Attribute Query

The last step in configuring an attribute source is to map values into the assertion to be sent in response to an attribute query on the Attribute Mapping Fulfillment tab.

Before you begin

For prerequisites and previous steps to configure the Attribute Query profile, see Configuring the Attribute Query profile in an SP connection on page 438.

Steps

- 1. For each attribute, select a source from the **Source** list and then choose or enter a value.
 - Context

When selected, the Value list populates with the available context of the transaction. Select the desired context from the list. At runtime, the context value is mapped to the value of the attribute in the SSO token.



Important:

If you are configuring an SP connection to bridge one or more identity providers to a service provider, consider mapping the original issuer of the assertions into an attribute by selecting Context as the source and Authenticating Authority as the value. This is important when bridging multiple identity providers to one service provider, where the service provider should take the information about the original issuer into consideration before granting access to protected resources.

For more information, see *Bridging multiple IdPs to an SP* on page 101.



Because the HTTP Request context value is retrieved as a Java object rather than text, use OGNL expressions to evaluate and return values (see Expression).

LDAP/JDBC/Other (when a datastore is used)

Values are returned from your datastore (if used). When you make this selection, the Value list is populated by the attributes from the datastore.

Expression (when enabled)

This option provides more complex mapping capabilities—for example, transforming incoming values into different formats (see Attribute mapping expressions on page 214). All of the variables available for text entries (see below) are also available for expressions.

No Mapping

Select this option to ignore the **Value** field, causing no value selection to be necessary.

Text

This can be text only, or you can mix text with references to any of the values from your userdatastore using this syntax:

```
${ds.attr-source-id.attribute}
```

where attr-source-id is the Attribute Source ID value (see Choosing a datastore for Attribute Query on page 440) and attribute is any of the datastore attributes you have selected.

There are a variety of reasons why you might hard code a text value. For example, if your SP's web application provides a service based on your company's name, you might provide that attribute value as a constant.

2. Click Next to save changes.

Defining issuance criteria for Attribute Query

Use the **Issuance Criteria** tab to define issuance criteria for the service provider (SP) Attribute Query profile.

About this task

On the Issuance Criteria tab, define the criteria to satisfy in order for PingFederate to further process a request. Use this token authorization feature to conditionally approve or reject requests based on individual attributes.

Begin this optional configuration by choosing the source that contains the attribute to verify. Some sources are common to almost all use cases, such as Mapped Attributes. Other sources depend on the type of configuration, such as JDBC. Irrelevant sources are automatically hidden. Once you select a source, choose the attribute to verify. Depending on the selected source, the available attributes or properties vary. Specify the comparison condition and the desired value to compare to.

You can define multiple criteria, which must all be satisfied in order for PingFederate to move a request to the next phase. A criterion is satisfied when the runtime value of the selected attribute matches or does not match the specified value, depending on the chosen comparison method. The multi-value contains ... or multi-value does not contain ... comparison methods are intended for attributes that can contain multiple values. Such a criterion is considered satisfied if one of the multiple values match or does not match the

specified value. Values are compared verbatim. If you require complex evaluations, including conditional criteria or partial matching, define them using attribute mapping expressions.



All criteria defined must be satisfied, or evaluated as true, for a request to move forward, regardless of how the criteria were defined. As soon as one criterion fails, PingFederate rejects the request and returns an error message.

Steps

1. From the **Source** list, select the attribute's source.

Depending on the selection, the Attribute Name list populates with associated attributes. See the following table for more information.

Source	Description
Context	Select to evaluate properties returned from the context of the transaction at runtime.
	Note:
	As the HTTP Request context value is retrieved as a Java object rather than text, attribute mapping expressions are more appropriate to evaluate and return values.
JDBC, LDAP, or other types of datastore	Select to evaluate attributes returned from a datastore, if configured.
Mapped Attributes	Select to evaluate the mapped attributes.

2. From the Attribute Name list, select the attribute to be evaluated.

3. From the **Condition** list, select the comparison method.

Available methods:

- equal to
- equal to (case insensitive)
- equal to DN
- not equal to
- not equal to (case insensitive)
- not equal to DN
- multi-value contains
- multi-value contains (case insensitive)
- multi-value contains DN
- multi-value does not contain
- multi-value does not contain (case insensitive)
- multi-value does not contain DN



Note:

The first six conditions are intended for single-value attributes. Use one of the multi-value ... conditions for PingFederate to validate whether one of the attribute values matches the specified value. When an attribute has multiple values, using a single-value condition causes the criteria to fail.

4. In the **Value** field, enter the comparison value.



Note:

Values are compared verbatim. If you require complex evaluations, including conditional criteria or partial matching, define them using attribute mapping expressions. For more information, see Attribute mapping expressions on page 214.

5. In the **Error Result** field, enter a custom error message.

The Error Result field is used by the StatusMessage element in the SAML response to the SP.

Using an error code in the Error Result field allows an application to process the code in a variety of ways; for example, display an error message or e-mail an administrator.

To use localized descriptions, enter a unique alias in the Error Result field, such as someIssuanceCriterionFailed. Insert the same alias with the desired localized text in the applicable language resource files, located in the <pf install>/pingfederate/server/ default/conf/language-packs directory.

If not defined, PingFederate returns ACCESS DENIED when the criterion fails at runtime.

- 6. Click Add.
- 7. Optional: Repeat to add more criteria.

- 8. If you require complex evaluations, including conditional criteria or partial matching, define them using attribute mapping expressions. For more information, see Attribute mapping expressions on page 214.
 - a. Click Show Advanced Criteria.
 - b. In the **Expression** field, enter the required expressions.
 - c. Optional: In the **Error Result** field, enter an error code or message.



Note:

If the expressions resolve to a string value instead of true or false, the returned value overrides the Error Result field value.

- d. Click Add.
- e. Optional: Click **Test**, enter values in the applicable fields, and verify the results.
- f. Optional: Repeat to add multiple criteria using attribute mapping expressions.

Specifying security policy

The Specify Security Policy tab allows you to specify the digital signing and encryption policy to which you and your partner have agreed.

About this task



Note:

The selections you make on this tab will trigger requirements for setting up Credentials. For more information, see Configuring credentials on page 445.

Steps

- 1. Select or clear the check boxes.
- 2. Click Next or Done.

Reviewing the Attribute Query configuration

Review and save your Attribute Query configuration changes on the **Summary** tab.

Steps

- To amend your configuration, click the corresponding tab title, then follow the configuration wizard to complete the task.
- To keep your changes, click **Done** and continue with the rest of the configuration.



When editing an existing configuration, you can also click Save as soon as the administrative console offers the opportunity to do so.

To discard your changes, click Cancel.

Configuring credentials

The Credentials tab provides the launching point for configuring security requirements you might need, depending on the federation protocol you are using and the choices you make.

Steps

To continue, click Configure Credentials.

See subsequent topics for configuration steps.

Configuring back-channel authentication (SAML)

Depending on your browser single sign-on (SSO) use cases, the administrative console prompts you to configure authentication requirements for inbound messages, outbound messages, or both.

About this task

See the following table for more information about the back-channel configuration (SAML) authentication requirements.

Use case	Back-channel authentication requirements	Back-channel messages
A connection is configured with a SAML ACS endpoint that uses the artifact binding on Protocol Settings # Assertion Consumer Service URL .	Inbound	Artifact resolution requests
A connection is configured with a SAML 2.0 SLO endpoint that uses the artifact binding on Protocol	Inbound	Artifact resolution requests
Settings # SLO Service URLs.		SOAP messages
A connection is configured with a SAML 2.0 SLO endpoint that uses the SOAP binding on Protocol Settings # SLO Service URLs .	Outbound	SOAP SLO messages
The SAML 2.0 Artifact binding is enabled on Protocol Settings # Allowable SAML Bindings.	Outbound	Outbound artifact resolution requests
The SOAP binding is enabled on Protocol Settings # Allowable SAML Bindings .	Inbound	Inbound SOAP messages
The SAML 2.0 Attribute Query profile is enabled on the Connection Options tab.	Inbound	Inbound Attribute Query requests

Steps

See subsequent topics for configuration steps.

Configuring authentication requirements for outbound messages

You can configure the authentication requirements used to validate outbound messages in PingFederate.

Steps

- 1. On the Back-Channel Authentication tab, in the Send to your partner section, click Configure.
- 2. On the **Outbound SOAP Authentication Type** tab, choose one or more authentication methods.

HTTP Basic

When selected, the administrative console prompts you to enter the credentials on the **Basic SOAP Authentication (Outbound)** tab.

You must obtain these credentials from your partner.

SSL Client Certificate

Applicable only if you specify an endpoint that uses HTTPS.

When selected, the administrative console prompts you to specify your client certificate on the **SSL Authentication Certificate** tab. If you have not yet created or imported the client certificate, click **Manage Certificates** to do so. For more information, see .



When exporting this client certificate for your partner, choose the Certificate Only option.

Digital Signature (Browser SSO profile only)

You select a signing certificate on the **Digital Signature Settings** tab.

This option leverages on the digital signature of the message.

Perform validation on partner's SSL server certificate when SSL used

By default, PingFederate validates your partner's HTTPS server certificate, verifying that the certificate chain is rooted by a trusted certificate authority (CA) and that the hostname matches the certificate's common name (CN).

Clear the associated check box if you do not want this validation to occur.

These options can be used in any combination or independently.

3. On the **Summary** tab, review your configuration and perform one of the following tasks.

Amend your configuration

Click the corresponding tab title and then follow the configuration wizard to complete the task.

Keep your changes

Click **Done** and continue with the rest of the configuration.



Tip:

When editing an existing configuration, you can also click Save as soon as the administrative console offers the opportunity to do so.

Discard your changes

Click Cancel.

Configuring authentication requirements for inbound messages

You can configure the authentication requirements used to validate inbound messages in PingFederate.

Steps

- 1. On the Back-Channel Authentication tab, in the Received from your partner section, click Configure.
- 2. On the **Inbound Authentication Type** tab, choose one or more authentication methods.

HTTP Basic

When selected, the administrative console prompts you to enter the credentials on the Basic SOAP Authentication (Inbound) tab.



Important:

If you are configuring more than one connection that uses the artifact or HTTP profile, you must ensure that the username is unique for each connection. You must communicate these credentials to your partner out-of-band.

SSL Client Certificate

When selected, the administrative console prompts you to specify the trust model and the related certificate settings on subsequent windows. See the next step.

Digital Signature (Browser SSO profile only)

You select a signing certificate on the **Signature Verification Settings** tab.

This option leverages on the digital signature of the message.

Require SSL

When selected, incoming HTTP transmissions must use a secure channel. This option is selected by default.

You can clear the check box if you do not require a secure channel and client certificate authentication.

For SAML 2.0, use these options in any combination or independently. For SAML 1.x, you must enable HTTP Basic authentication, client certificate authentication, or both. You can also add digital signing to ensure message integrity.

3. If you chose SSL Client Certificate in the previous step, select a trust model on the Certificate Verification Method tab.

Anchored

The partner certificate must be signed by a trusted certificate authority (CA). Optionally, you can also restrict the issuer to a specific Trusted CA to mitigate potential man-in-the-middle attacks and to provide a means to isolate certificates used by different connections. The CA's certificate must be imported into the PingFederate Trusted CA store on the Trusted CAswindow...

Unanchored

The partner certificate is self-signed or you want to trust a specified certificate.



Note:

When anchored certificates are used between partners, certificates can be changed without sending the update to your partner. If the certificate is unanchored, any changes must be promulgated.

For more information, see .

Trust model	Subsequent steps	
Anchored	On the Subject DN tab:	
	 a. Enter the Subject DN of the certificate. b. Optionally, select the Restrict Issuer check box and enter the Issuer DN of the certificate. 	
	Important:	
	Consider enabling this option to mitigate potential man-in-the-middle attacks and to provide a means to isolate certificates used by different connections.	
Unanchored	On the SSL Verification Certificate tab, select the client certification from your partner.	
	If you have not yet imported the client certificate from your partner, click Manage Certificates to do so. For more information, see .	

4. On the **Summary** tab, review your configuration and perform one of the following tasks.

Amend your configuration

Click the corresponding tab title and then follow the configuration wizard to complete the task.

Keep your changes

Click **Done** and continue with the rest of the configuration.



Tip:

When editing an existing configuration, you can also click Save as soon as the administrative console offers the opportunity to do so.

Discard your changes

Click Cancel.

Configuring digital signature settings

Configure digital signing in the PingFederate administrative console. Digital signing is required for browserbased single sign-on (SSO) tokens and single logout (SLO) messages sent through POST or redirect bindings.

About this task

Digital signing is also required for WS-Trust STS service provider (SP) connections, for signing the outbound SAML security tokens.

For browser-based SSO, digital signing is not always required for profiles using the artifact or SOAP bindings unless you chose to sign the SAML assertion on Protocol Settings # Signature Policy, or the artifact resolution messages on Back-Channel Authentication # Outbound SOAP Authentication Type.

If digital signing is not required, PingFederate does not show the Digital Signature Settings tab.

Steps

- 1. On the **Digital Signature Settings** tab, select the certificate that you will use to sign the SSO tokens and SLO messages for the SP.
- 2. Select a signing certificate from the Signing Certificate list.

If you have not yet created or imported your certificate into PingFederate, click Manage Certificates. For more information, see Manage digital signing certificates and decryption keys on page 608.



Note:

For WS-Federation connections using JSON Web Tokens (JWTs), only EC and RSA certificates are supported. RSA certificates must have a minimum key size of 2,048 bits. The Signing Certificate list automatically filters out certificates that do not meet these requirements.

3. Optional: Select the Include the certificate in the signature <KeyInfo> element check box if you have agreed to send your public key with the message.



Note:

For WS-Trust STS, the <KeyInfo> element in the SAML token includes a reference to the certificate rather than the full certificate by default unless this check box is checked.



Note:

This step is not applicable to WS-Federation connections using JWTs.

Select the Include the raw key in the signature <KeyValue> element check box if your partner agreement requires it.

4. Optional: Select the signing algorithm from the list.

The default selection is RSA SHA256 or ECDSA SHA256, depending on the Key Algorithm value of the selected digital signing certificate. Make a different selection if you and your partner have agreed to use a stronger algorithm.

Configuring signature verification settings (SAML 2.0)

You can configure the signature verification settings for the certificates in the PingFederate administrative console.

About this task

Depending on your partner agreement, digital signature processing might be required.

If you choose to require digital signatures on SAML 2.0 authentication requests on Protocol Settings # Signature Policy or inbound messages on Back-Channel Authentication # Inbound Authentication Type, you must configure the required certificate information that PingFederate can use to verify the signed messages.

The **Signature Verification Settings** tab is the launching point for this task. If digital signature verification is not required, the **Signature Verification Settings** tab is not shown.

Steps

- 1. On the Signature Verification Settings tab, click Manage Signature Verification Settings.
- 2. On the **Trust Model** window, select a trust model on the **Certificate Verification Method** tab.

Anchored

The partner certificate must be signed by a trusted certificate authority (CA). Optionally, you can also restrict the issuer to a specific Trusted CA to mitigate potential man-in-the-middle attacks and provide a means to isolate certificates used by different connections. The CA's certificate must be imported into the PingFederate Trusted CA store on Security # Certificate & Key Management # Trusted CAs.



Important:

If you are using the redirect binding for single logout (SLO), you cannot use anchored certificates because SAML 2.0 does not permit certificates to be included using this transport method.

Unanchored



When anchored certificates are used between partners, certificates can be changed without sending the update to your partner. If the certificate is unanchored, any changes must be promulgated.

For more information, see *Digital signing policy coordination* on page 88.

Trust model	Subsequent steps
Anchored	On the Subject DN tab:
	 a. Enter the subject DN of the certificate or extract it from your service provider (SP) partner's certificate if the certificate is stored on an accessible file system. b. Optionally, select the Restrict Issuer check box and enter the Issuer DN of the certificate. Alternatively, extract it from your partner's certificate.
	Important:
	You can enable this option to mitigate potential man-in-the-middle attacks and to provide a means to isolate certificates used by different connections.
Unanchored	On the Signature Verification Certificate tab:
	a. Select a certificate from the list.
	If you have not yet imported the certificate from your partner, click Manage Certificates to do so. For more information, see Managing certificates from partners on page 622. b. Optionally, select additional certificates.
	Note:
	When configured, PingFederate considers a digital signature valid so long as it can verify the signature using one of the certificates from this list.
	Tip:
	This is useful in situations where your partner has sent you a certificate to replace the current certificate. Adding this second certificate allows PingFederate to continue validating digital signatures as the partner switches to the new signing certificate.
	It also adds support for scenarios where your partner uses a pool for certificates to sign its messages. Adding these certificates ensures digital signatures can be validated as the partner rotates its signing certificates.

3. On the **Summary** tab, review your configuration and perform one of the following tasks.

Amend your configuration

Click the corresponding tab title and then follow the configuration wizard to complete the task.

Keep your changes

Click **Done** and continue with the rest of the configuration.



When editing an existing configuration, you can also click **Save** as soon as the administrative console offers the opportunity to do so.

Discard your changes

Click Cancel.

Selecting an encryption certificate

For browser-based single sign-on (SSO), if you choose to encrypt all or part of an SSO assertion on Protocol Settings # Encryption Policy, you must identify the certificate that PingFederate can use to do SO.

About this task

You must also select a certificate if your requirements include encrypting an assertion in response to an attribute query on Attribute Query # Security Policy.

For WS-Trust security token service (STS), this configuration is also required if you enabled the Generate Key for SAML Holder of Key Subject Confirmation Method or Encrypt SAML 2.0 Assertion option, or both, on WS-Trust # Protocol Settings.

If encryption is not required, the **Select XML Encryption Certificate** tab is not shown.

Steps

1. Optional: Select an option under **Block Encryption Algorithm**.



Important:

Due to the import restrictions of some countries, Oracle Server Java SE Runtime Environment (JRE) 8 has built-in restrictions on available cryptographic strength (key size). To use larger key sizes, enable the Java Cryptography Extension (JCE) unlimited strength jurisdiction policy. For more information, see the Java 8 release notes in Oracle's documentation.

For Oracle Java SE Development Kit 11, the JCE jurisdiction policy defaults to unlimited strength. For more information, see the Oracle JDK Migration Guide in Oracle's documentation.

The default selection is AES-128.

For more information about XML block encryption and key transport algorithms, see XML Encryption Syntax and Processing from W3C.

2. Select an option under Key Transport Algorithm.



Note:

Due to security risks associated with the RSA-v1.5 algorithm used for key transport, it is no longer available for new connections. Existing connections in which this algorithm is configured continue to support it. However, you should upgrade such connections to use the newer algorithm RSA-OAEP.

The default selection is **RSA-OAEP**.

3. Select a partner certificate from the list.

If you have not imported the certificate from your partner, click Manage Certificates to do so. For more information see .

Selecting a decryption key (SAML 2.0)

To enable inbound encryption in PingFederate, you must select a certificate on the decryption key.

About this task

When you choose to encrypt the name identifier (SAML SUBJECT) on Protocol Settings # Encryption **Policy**, you can also allow the service provider (SP) to encrypt the name identifier in its single logout (SLO) requests, if the SP-initiated single sign-on (SSO) profile is enabled for the connection. To enable this inbound encryption, you must specify at least one certificate on the Select Decryption Keys tab.

If decryption is not required, the **Select Decryption Keys** window is not shown.

Steps

- 1. Select the primary XML decryption key from the list.
 - If you have not created or imported your certificate into PingFederate, click Manage Certificates. For more information, see .
- 2. Optional: Select the secondary XML decryption key from the list.

Reviewing SP credential settings

You can review, modify, save, and discard changes to your service provider (SP) credential settings in the PingFederate administrative console.

Steps

- To amend your configuration, click the corresponding tab title, then follow the configuration wizard to complete the task.
- To keep your changes, click **Done** and continue with the rest of the configuration.



When editing an existing configuration, you can also click Save as soon as the administrative console offers the opportunity to do so.

• To discard your changes, click Cancel.

Configuring outbound provisioning

PingFederate's outbound provisioning allows an identity provider (IdP) to create and maintain user accounts at standards-based partner sites using System for Cross-domain Identity Management (SCIM) as well as select-proprietary provisioning partner sites that are protocol-enabled.

About this task

For more information, see *Outbound provisioning for IdPs* on page 97.

Note:

This configuration task is presented in the administrative console only when you enable the Outbound **Provisioning** protocol. For more information, see *Choosing an SP connection type* on page 403.

Steps

Go to Applications # Integration # SP Connection. To continue, click Configure Provisioning.



Defining a provisioning target

You can define a provisioning target, including the provider's web-service endpoint for provisioning users and, if required, credentials that PingFederate uses for authentication to the provisioning API for the service provider (SP).

About this task





Note:

The target configuration settings vary among System for Cross-domain Identity Management (SCIM) outbound provisioning and various software as a service (SaaS) provisioning.

For SCIM provisioning to PingOne for Enterprise, sign on to the *PingOne admin portal* and review the target information on the Setup # Identity Repository tab.

For any SaaS Connector target, see documentation in the add-on distribution package.

The following steps describe the fields required for the bundled PingFederate provisioning plugin for SCIM partners.

Steps

- Enter the endpoint for managing users in the Users Resource URL field, such as, https:// example.com/v1/Users.
 - This field is always required for SCIM outbound provisioning.
- 2. Go to Applications # Integration # SP Connections # SP Connection # Configure Channels.

 ${\it 3.} \ \ {\it On the \ \textbf{Custom SCIM Attributes}} \ tab, \ configure \ the \ rest \ of \ the \ outbound \ provisioning \ settings.$

See the following table for detailed information about each field.

Field	Description
Groups Resource URL	The partner's group management endpoint. For example, https://example.com/v1/Groups.
	Required if the partner supports this notion and groups should be provisioned.
Authentication Method	The authentication scheme that the partner's endpoints support.
	Available options:
	 None Basic Authentication (Default) OAuth 2.0 Bearer Token - Uses the resource owner grant type by submitting the client ID, client secret, username, and password to the configured token endpoint URL in exchange for an access token that will be sent in each SCIM request.
User, and	Valid credentials to access the partner's endpoint.
Password	Required if Basic Authentication is the selected authentication method.
Client ID, Client Secret, and	Valid OAuth client credentials and token endpoint to access the partner's endpoint.
Token Endpoint URL	Required if OAuth 2.0 Bearer Token is the selected authentication method.
SCIM SP Supports	Clear this check box if the partner does not support PATCH updates.
Patch Updates	For information about PATCH, see the SCIM specification (www.simplecloud.info/specs/draft-scim-api-01.html#edit-resource-with-patch).
	This check box is selected by default.
Provision Groups with Distinguished Name	Select this check box to provision groups by supplying complete LDAP distinguished names (DNs), rather than only common names (CNs), to identify groups.
	Some SCIM partners, including PingOne for Enterprise, allow administrators to parse full DNs when necessary, such as in the case of duplicate CNs, to determine group access mapping to specific applications based on other DN elements. Consult the partner for its requirement.
	This check box is selected by default.

4. Click Next.



Note:

For some provisioning plugins, including the built-in SCIM outbound provisioner, when you enter or change credentials and click Next, PingFederate immediately tests connectivity to the target.

Specifying custom SCIM attributes

You can configure simple, multivalued, and complex custom System for Cross-domain Identity Management (SCIM) attributes in PingFederate.

About this task

PingFederate supports SCIM attributes in the core schema and custom attributes through a schema extension.



Note:

Custom attributes are optional. If your use case does not require any additional attributes, click Next on the Custom SCIM Attributes tab.

To support custom attributes, you must specify the schema extension and the custom attributes in the connection. There are four attribute types:

- Simple attributes
- Simple multivalued attributes
- Complex attributes
- Complex multivalued attributes

The following fragment illustrates a SCIM message supporting schema extension

urn:scim:schemas:extension:custom:1.0 with four attributes, one of each attribute type. The table afterward describes the details of each attribute.

```
"userName": "CBrown",
"active":true,
"schemas":[
  "urn:scim:schemas:core:1.0",
  "urn:scim:schemas:extension:custom:1.0"
],
"urn:scim:schemas:extension:custom:1.0":{
  "supervisor": "JSmith",
  "territories":[
    "Montana",
     "Idaho",
     "Wyoming"
  "options":{
     "quantity":"10000",
"strike" :"5.25",
"first" :"2017-12-01",
"last" :"2025-03-31"
  },
  "tablets":[
     {
       "model" :"8086",
       "serial":"5500-2020-965",
"type" :"office"
     },
        "model" : "8088",
       "serial":"5500-2040-151",
"type" :"remote"
  ]
}
```

Attribute Name	Attribute Type	Sub-Attributes (Complex)
supervisor	Simple	Not applicable
territories	Simple multivalued	Not applicable
options	Complex	quantity, strike, first, and last
tablets	Complex multivalued	model, serial, and type.
		Note: type is a reserved sub-attribute for a complex multivalued attribute.



For more information about SCIM and attribute types, see the website www.simplecloud.info.

1. Go to Applications # Integration # SP Connection # Configure Channels. Specify the URI of the schema extension in the Extension Namespace field.





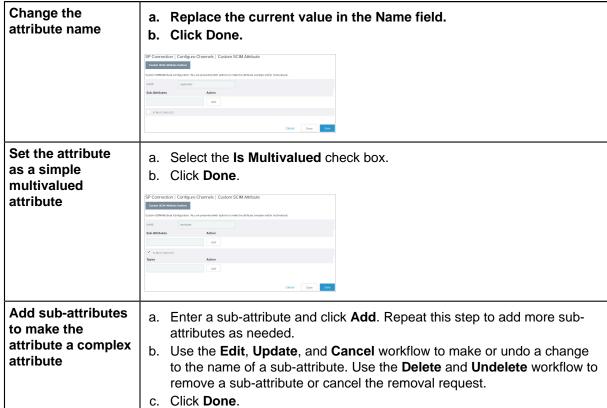
The default value is urn:scim:schemas:extension:custom:1.0. You can keep this value if your partner identifies custom attributes by this URI in its SCIM messages.

2. Enter an attribute name and click **Add** to add a custom attribute. Repeat this step to add more custom attributes as needed.



Use the **Delete** and **Undelete** workflow to remove or cancel the removal request of existing custom attributes.

3. Click **Edit** next to the custom attribute to perform one of the following tasks.



Managing channels

You can manage provisioning channels in the PingFederate administrative console to add, modify, and delete channels.

About this task

A provisioning channel is a mapping configuration between user attributes contained in a source user store and attributes supported or required by the targeted software-service application. You can have multiple channels to the same target as needed, such as if your organization has separate LDAP stores, or different

nodes in the same store, for various user groups needing single sign-on (SSO) access and provisioning to the same domain.



There can be only one provisioning target per connection. If your organization subscribes to multiple domains for which you need provisioning support, you need a separate service provider (SP) connection for each domain.



Go to Applications # Integration # SP Connection # Configure Channels. On the Manage Channels tab, you can perform the following tasks:

Steps

- To add a new channel, click Create.
 - Alternatively, you can create a new channel by copying an existing channel and making other required changes.
- To modify existing channel settings, select an existing channel.
- To remove or cancel the removal request of the existing channel, use the **Delete** and **Undelete** workflow.

Specifying channel information

When configuring channels in PingFederate, you can specify channel information such as the channel name, max threads, and timeout.

About this task

On the Channel Info tab, specify a unique identifier for the channel and adjust the values for the Max Threads and Timeout fields as needed to optimize data-transfer performance, particularly if large numbers of records need to be provisioned at the target site.



Steps

- 1. Go to Applications # Integration # SP Connection # Configure Channels # Channel. In the Channel Name field, enter a channel name.
 - If you are copying a channel, you must enter a new value in this field.
- 2. Optional: Update the values for the Max Threads and Timeout fields.
 - The **Timeout** value applies only when the **Max Threads** value allows multiple threads.
- 3. Click Next.

Identifying the source datastore

You can identify the source datastore for the service provider (SP) channel configuration in the PingFederate administrative console.

About this task

PingFederate supports PingDirectory, Microsoft Active Directory, Oracle Unified Directory, and Oracle Directory Server as source user repositories for outbound provisioning. However, you can use other types of LDAP servers, either identifying them as **Generic** or registering them with PingFederate. For more information, see the sample.template.txt in the $\protect\protec$

Information from your user-datastore is used to supply mapped values for each user attribute required by the service provider (SP).



Steps

- 1. Go to Applications # Integration # SP Connections # SP Connection # Configure Channels # Channel.
- 2. On the **Source** tab, choose the LDAP store to use for this channel.

If the datastore you want is not shown in the list, PingFederate is not configured to access the store. To create a connection to the datastore, click **Manage Data Stores**.

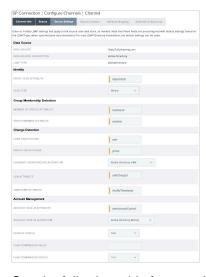
3. Click Next.

Modifying source settings

You can modify the source settings for the datastore configuration in the PingFederate administrative console. You can add, change, and remove user information.

About this task

The **Source Settings** tab shows the default configuration of the datastore selected on the **Source** tab, including settings used by the PingFederate provisioner to determine when user information is added, changed, or removed.



See the following table for more information about each field.

Field	Description
Entry GUID Attribute	The name of the attribute in the datastore representing the user's GUID.
GUID Type	Indicates whether the GUID is stored in binary or text format. Active Directory is always binary. Other LDAP stores most often use text.
Member of Group Attribute	A multivalued user attribute containing the distinguished names (DNs) of the groups to which an entry belongs. This attribute does not apply to some LDAP servers, including Oracle Unified Directory and Oracle Directory Server. The attribute in the next row is used instead. PingDirectory and Microsoft Active Directory use both values to provide a two-way mapping between user and group objects.
Group Member Attribute	The name of a multivalued group attribute used to track membership in the group using either DN or GUID values.
User objectClass	The LDAP object class to which user entries belong, used to restrict search results to user entries only.
Group objectClass	The LDAP object class to which group entries belong, used to restrict search results to group entries only.
Changed Users/Groups Algorithm	The method by which PingFederate determines if user records have been updated or new records added, thus requiring provisioning updates at the target site. The three choices are:
	 Active Directory USN – For Active Directory only, this algorithm queries for update sequence numbers on user records that are larger than the last time records were checked. Timestamp – Queries for timestamps on user records that are not older than the last time records were checked. This check is more efficient from the perspective of the PingFederate provisioner but can be more time consuming on the LDAP side, particularly with Oracle Unified Directory and Oracle Directory Server. Timestamp No Negation – Queries for timestamps on user records that are newer than the last time records were checked. This algorithm is recommended for Oracle Unified Directory and Oracle Directory Server.
USN Attribute	The name of the attribute used to store the update sequence number. Applicable when the Active Directory algorithm is chosen in the row above.
Timestamp Attribute	The name of the attribute used to store the timestamp on user records.
	Note: This attribute name is case-sensitive. Ensure the attribute name matches the name your directory uses. For example, in PingDirectory and Oracle, the attribute is modifyTimestamp, but in Microsoft Active Directory, the attribute is modifyTimeStamp. They have different capitalization.
Account Status Attribute	The name of the attribute in which the user's account status, active or inactive, is stored. For example, Active Directory = userAccountControl and Oracle Directory Server = nsaccountlock.

If you are using PingDirectory, Microsoft Active Directory, Oracle Unified Directory, or Oracle Directory Server, in most cases no changes are needed on this tab unless your datastore uses a customized schema.

If you are using a different LDAP directory, you must supply the required information on this tab unless you have defined a template for the datastore. For more information, see the sample.template.txt in the install/pingfederate/server/default/conf/template/ldap-templates directory.

Steps

- 1. Modify the settings, as needed.
- 2. Click Next.

Specifying a source location

You can indicate on the **Source Location** tab where PingFederate should look for user records in the datastore.

About this task

The same location can be used to retrieve user-group distinguished names (DNs) for maintaining corresponding groups at the service provider (SP).





Note:

Groups provisioning is supported for System for Cross-domain Identity Management (SCIM) and the Google Apps Connector (version 2.0 and higher) but might not be supported for other software as a service (SaaS) Connectors. If not, the associated fields under Groups on the Source Location tab are inactive. Support for the feature might become available in future SaaS Connector releases. See the documentation in your add-on distribution package.

Steps

1. Go to Applications # Integration # SP Connections # Configure Channels # Channel. In the **Base DN** field, enter the base DN where user records are stored.

PingFederate looks only at this node level, or below it, for user accounts and groups (when applicable) that need to be provisioned based on the conditions set in the next step.

2. Specify group membership information or an LDAP filter to search for users, and groups when applicable, to be provisioned. For more information, see the following table.

Object	Field description
Users	Group DN
	The distinguished name (DN) of a group in the user repository whose member groups should be provisioned.
	Optionally, select the Nested Search check box to include users that are members of the specified group through nested group membership. Nested group membership is preserved for SCIM provisioning, and SaaS provisioning if the vendor and the SaaS Connectors support hierarchical structure in groups.
	Note:
	The Nested Search feature is available when Microsoft Active Directory, Oracle Unified Directory, or Oracle Directory Server is selected as the source user repository. For more information, see <i>Identifying the source datastore</i> on page 461.
	Filter
	An LDAP search filter that returns user objects representing the users that should be provisioned.
	For information about LDAP filters, see your LDAP documentation. You might need to escape any special characters.
	Important:
	The Group DN field is ignored when a Filter field value is configured.

ned.
s that ership. nd SaaS urchical
rectory, s the source
groups
ation. You
ıred.
s will be

3. Click Next.

Mapping attributes

Mapping attributes determines how attributes from your user store are mapped to the System for Crossdomain Identity Management (SCIM) attributes in the core schema and custom attributes through a schema extension or to the provisioning fields supported for your organization's software as a service (SaaS) customer account.

About this task

Edit the mapping of attributes from the local datastore into fields specified by the service provider (SP).





Important:

If you are provisioning for SCIM, your SP can make one or more optional core attributes mandatory. For more information, see the SCIM documentation from the SP or the SCIM Resource Schema representation.

Tip:

For non-SCIM SaaS connectors, PingFederate automatically retrieves from the vendor the Field Names shown on this tab, but only on the first pass through the configuration flow. If you are using this configuration to modify an existing mapping configuration, click Refresh Fields to synchronize the list with the target if needed.

For each field, the Specify Attribute Mapping option provides a means of adding or modifying the mapping details.



Note:

All required attributes listed in the **Field Name** column, indicated with asterisks, must be mapped. Click View Partner Field Specifications for a summary of requirements for all fields specified for the target partner.

For some fields, PingFederate preselects LDAP attributes commonly used to store the required values.

Steps

- 1. Go to Applications # Integration # SP Connections to open the SP Connections configuration window.
- 2. To edit an existing SP connection, open an SP connection by clicking on its name in the Connection Name column.
- 3. On the Outbound Provisioning tab, click Configure Provisioning to open the Configure Channels configuration window.



Note:

The Outbound Provisioning tab is only visible after you select the Outbound Provisioning check box and from the **Type** list, select the type on the **Connection Type** tab.

- 4. Go to the Manage Channels tab.
- 5. Select a channel.
- 6. Go to the Attribute Mapping tab.
- 7. To edit a field, click **Edit** under **Action**.



If you have specified any custom attributes, they are listed at the end of the Attribute Mapping configuration.

8. On the Attribute Mapping tab, provide mapping details.



For most fields, if you map more than one attribute from your datastore into a single field at the target location, then you must use an OGNL expression to indicate how to combine the attribute values.

The only exception is the LDAP Attributes Map field, which is provided primarily to support SCIM attributes specific to PingOne for Enterprise. This field can contain multiple attributes without using OGNL.

10. Click Next.

Specifying mapping details

Define specific mapping information for each field, required or optional, for provisioning as needed.



If end-users at your site are permitted to edit some of their own attributes directly in the LDAP store, ensure that the attributes are restricted and do not include any needed by the service provider to grant permissions.

Defining mapping information for a standard attribute

Before you begin

- Go to Applications # Integration # SP Connections to open the SP Connections configuration
- To edit an existing SP Connection, open an SP Connection by clicking on its name in the Connection Name column.
- On the Outbound Provisioning tab, click Configure Provisioning to open the Configure Channels configuration window.



Note:

The Outbound Provisioning tab, is only visible after you select in the OutBound Provisioning check box and from the Type list, select the type on the Connection Type tab.

- Go to the Manage Channels tab.
- Click Edit in the Action column to edit a field.



If you have specified any custom attributes, they are listed at the end of the Attribute Mapping configuration.

Steps

1. On the Specify Attribute Mapping tab, provide mapping details.

2. Select the class containing a user-store attribute under Root Object Class that you want to map to the provisioning attribute shown under **Field Name**.



Note:

For some fields, you might not need to map specific user attributes. If so, supply a value in the **Default** Value field instead, skip this step, and go to step 5. For certain attributes, you can specify both LDAP attributes and a default value, as needed.

- 3. Select the source attribute from the class under Attribute. Click Add Attribute.
- 4. Repeat the previous steps to add additional applicable attributes to use in a mapping expression.



You must add an attribute for it to be used in an expression.

5. Optional: If one or more attributes are specified, enter or select a default value from the Value **Definition** list.

A list appears for this field if the vendor requires a choice among specified values. When an expression is also supplied, the default value is sent during provisioning if an error occurs when evaluating the expression.

- 6. If more than one attribute is used for mapping fields other than LDAP Attributes Map, enter an expression.
 - a. Click **Edit** to create and validate the expression.
- 7. Select one or more processing options.

Processing option	Description
Create Only	The field is provisioned only once and not subsequently updated.
	Note:
	For SCIM, the Password attribute should be passed only when creating a user or updating the password. Select Create Only to limit when the Password attribute is passed.
Trim	Removes any white space from the attribute values.
Mask Log Values	Determines whether sensitive information, such as the Password attribute, will be masked in PingFederate log files.
Upper Case, Lower Case, or None	Transforms the attribute values to the case indicated unless the default, None option, is selected.
Parsing > Extract CN from DN	For attributes in the form of a distinguished name (DN), such as Group DNs in Active Directory, maps only the common name portion of the DN.
Parsing > Extract Username from Email	For attributes containing an email address, maps only the username.

Defining mapping information for a custom attribute

Steps

1. Select a sub-attribute from the Attribute ID list.



Note:

Applicable only to complex attributes or complex multivalued attributes, see Specifying custom SCIM attributes on page 456.

2. Select the class containing a user-store attribute under Root Object Class that you want to map to the provisioning attribute shown under Field Name.



Note:

For some fields, you might not need to map specific user attributes. If so, supply a value in the **Default** Value field instead, skip this step, and go to step 5.For certain attributes, you can specify both LDAP attributes and a default value, as needed.

- 3. Select the source attribute from the class under LDAP Attribute. Click Add Attribute.
- 4. Select one or more processing options.

Processing option	Description
Create Only	The field is provisioned only once and not subsequently updated.
	For SCIM, the Password attribute should be passed only when creating a user or updating the password. Select Create Only to limit when the Password attribute is passed.
Trim	Removes any white space from the attribute values.
Mask Log Values	Determines whether sensitive information, such as the Password attribute, will be masked in PingFederate log files.
Upper Case, Lower Case, or None	Transforms the attribute values to the case indicated unless the None option is selected, the default.
Parsing > Extract CN from DN	For attributes in the form of a distinguished name (DN), such as Group DNs in Active Directory, maps only the common name portion of the DN.
Parsing > Extract Username from Email	For attributes containing an email address, maps only the username.

5. Enter a default value.

6. Click Add Mapping.



Note:

For complex attributes or complex multivalued attributes, repeat these steps to map additional subattributes as needed.

7. Click Done.

Reviewing channel settings

When you finish setting up a channel, you can choose to activate it immediately or activate the channel as needed.

About this task

On the **Activation & Summary** tab, review, activate, deactivate, or save your channel settings.



A service provider connection must be active for any provisioning channels to be enabled. You can deactivate a channel at any time. When a channel is inactive, provisioning is suspended but single sign-on (SSO) and single logout (SLO) transactions can still occur if an associated connection is active.

Steps

- 1. Go to Applications # Integration # SP Connections # SP Connection # Configure Channels # Channel.
- 2. On the Activation & Summary tab, select Active or Inactive to toggle the status. Click Save.



CAUTION:

When a channel is activated, initial provisioning occurs as soon as the synchronization-frequency time period expires. See Configuring outbound provisioning settings on page 839. The default is 60 seconds. Because initial provisioning can consume considerable processing time, depending on the amount of data that needs to be transmitted, administrators should plan accordingly.

3. To modify channel settings, click the associated heading in the **Summary** column.



Important:

Click **Save** to save the channel configuration.

Reviewing SP connection settings

When you finish creating or modifying a service provider (SP) connection, you can review the connection settings and toggle the connection status.

About this task

On the **Activation & Summary** tab, review, amend, discard, or save your changes.

Steps

- To amend your configuration, click the corresponding tab title and then follow the steps to complete the task.
- To keep your changes, click Save.

To discard your changes, click Cancel.

Result



Important:

When creating a new connection, the default connection status is **Enabled** when you reach the **Activation** & Summary tab.

Whether you choose to disable a new connection now or later, you must click Save on the Activation & **Summary** tab if you want to keep the new connection.

The SSO Application Endpoint provides a sample URL at the /idp/startSSO.ping application endpoint that webmasters or web application developers at your site can use to invoke single sign-on for the connection. For a list of supported parameters, see Viewing IdP application endpoints on page 396.

SP affiliations

A service provider (SP) affiliation is a SAML 2.0 specification that permits a group of service providers to make use of the same persistent name identifier for account linking.

SP affiliations are useful when multiple SPs share a business relationship in which users need services from each affiliated provider. By agreement among the affiliation members, the same pseudonym can be used to populate the SAML SUBJECT of assertions sent to all of the SP partners contained in this affiliation.



Note:

Each connection in the affiliation must be configured to use the same identity provider adapter instance for generating account links. For more information, see Managing authentication source mappings on page 417.

Managing SP affiliations

Use service provider (SP) affiliations when multiple SP's share a business relationship in which users need services from each affiliated provider. To support these needs, edit and delete your SP affiliations in PingFederate as needed.

About this task

In **SP Affiliations** window, create, edit, or delete an SP affiliation.

Steps

 To create a new SP affiliation, go to System # Protocol Metadata # SP Affiliations to open the SP Affiliations window. Click Create Affiliation.

Result: This is will open the **Create an Affiliation** window configuration.

- To edit or delete an SP affiliation, choose from the following options. Choose from:
 - To edit an SP affiliation, on the SP Affiliations window, select the affiliation by its ID and follow the configuration wizard to complete the task.
 - To delete an SP affiliation, in the SP Affiliations window, click Delete under Action for the SP affiliation.

Importing affiliation metadata

You can import a metadata file for an identity provider (IdP) to send containing information that automatically specifies members of a service provider (SP) affiliation, describing this affiliation, to provide authentication for account linking.

About this task

On the Import Metadata tab, import a metadata file, or click Next if you do not want to import a file.

Steps

 Go to System # Protocol Metadata # SP Affiliations to open the SP Affiliations window. Click Create Affiliation.

Result: This is will open the **Create an Affiliation** window configuration.

- To import a metadata file or not and move onto the next step, choose from the following options.
 Choose from:
 - To import a metadata file, on the Import Metadata tab, click Choose File to upload it. Click Next.
 - If you do not have a metadata file, on the **Import Metadata** tab, click **Next**.

Entering affiliation information

If you did not import a metadata file to describe a new service provider (SP) affiliation, enter the information manually to identify your SP affiliation to support authentication for account linking.

About this task

On the Affiliation General Info tab, enter the affiliation ID information, as described in the following table.

Steps

 To create a new SP affiliation, go to System # Protocol Metadata # SP Affiliations to open the SP Affiliations window. Click Create Affiliation.

Result: This is will open the **Create an Affiliation** window configuration.

• On the **Affliation General Info** tab, enter the following information.

Affiliation ID fields and descriptions

Field	Description
Affiliation ID	A unique identifier for this affiliation. This value serves as the Name ID qualifier for SAML assertions sent to affiliated SP partners.
Affiliation Owner	Any SAML 2.0 SP connection can serve as the Owner.

If you imported a metadata file, this information is already supplied. However, you can change the Affiliation ID or select a different Affiliation Owner, if required.

When finished, click Next.

Managing affiliation membership

Manage the list of service provider (SP) connections that are part of this affiliation. Configure each of the SP connections in an affiliation to generate opaque pseudonyms using the same adapter attributes.

About this task

On the **Affiliation Membership** tab, create and manage a list of SP connections to be included in the affiliation.

If you imported a metadata file, this information is already supplied. However, you can add or remove connections from the affiliation.

 Go to System # Protocol Metadata # SP Affiliations to open the SP Affiliations window. Click Create Affiliation.

Result: This is will open a new **Create an Affiliation** configuration window.

- Click on the Affliation Membership tab.
- To add an SP partner connection to the affiliation, from the SP Connection Name list, select the connection. Click Add.



Important:

Each connection in the affiliation must be configured to use the same identity provider (IdP) adapter instance for generating account links. For more information, see Managing authentication source mappings on page 417.

To remove a member of the affiliation, click **Delete** under **Action** for the connection.



If you delete an affiliation member supplied by an imported metadata file and then save the affiliation, that connection will not appear in the drop-down list for re-adding in the future.

When finished, click Next.

Reviewing an SP affiliation

Review the summary information for a service provider (SP) affiliation and make changes or save as needed.

About this task

On the Activation & Summary, review, amend, discard, or save your changes.

Steps

- To amend your configuration, click the corresponding tab title and then follow the steps to complete the task.
- To keep your changes, click Save.
- To discard your changes, click Cancel.

OAuth configuration

OAuth is an open standard for authorization. To use PingFederate as an OAuth authorization server (AS), configure the OAuth AS settings as described in this section.



Service providers can also add OAuth capabilities to the Browser SSO configuration for identity provider (IdP) connection partners, see Configure OAuth attribute mapping.

Administrators can configure PingFederate to support the OAuth grant types that applications require.

Steps

- 1. To configure the authorization server settings, go to System # OAuth Settings # Authorization Server Settings. For more information, see Configuring authorization server settings.
- 2. Define any number of optional common scopes and exclusive scopes, create scope groups from optional scopes as needed, and enter an optional description for the default scope in the System # OAuth Settings # Scope Management window.
- 3. Create one or more access token management instances in the Applications # OAuth # Access Token Management window.



You can also define the access token attribute contract for an access token management instance in this window.

4. Configure one or more entries to map attributes from authentication sources to the persistent grants.

Authorization Code or Implicit

- Map attributes from an identity provider (IdP) adapter instance to the persistent grants in Authentication # OAuth # IdP Adapter Grant Mapping.
- Map attributes from an IdP connection to the persistent grants in IdP Connection # **Browser SSO # OAuth Attribute Mapping.**
- Create an authentication policy contract (APC) using the Policy Contracts window, define an authentication policy to map attributes from the authentication sources (IdP adapter instances, IdP connections, or both) to the APC, and map attributes from the APC to the persistent grants using the Authentication Policy Contract Grant Mapping window.



If you are using a combination of authentication policies, APCs, and APC mappings, you can skip the IdP Adapter Grant Mapping and OAuth Attribute Mapping configurations.

Resource Owner Password Credentials

 Map attributes from a password credential validator instance to the persistent grants using the Authentication # OAuth # Resource Owner Credentials Grant Mapping configuration wizard.



Note:

This is the first stage of the two-stage access token mapping process through the persistent grants.

5. Configure one or more entries to map attributes from the persistent grants (or the authentication sources directly) to the attribute contract of your access token management instances in the

Applications # OAuth # Access Token Mapping window. Additionally, you can configure a mapping for clients using the client credential grant type.



Note:

This is the second stage of the two-stage access token mapping process through the persistent grants. For more information about the access token mapping process, see *Mapping OAuth attributes* on page 81.

- 6. For the client-initiated backchannel authentication (CIBA) flow, configure one or more CIBA authenticator instances and then one or more CIBA request policies.
- 7. For the JSON web token (JWT) Bearer or SAML 2.0 Bear assertion grants flow, configure a mapping in IdP Connection # OAuth Assertion Grant Attribute Mapping.



This use case exchanges a JWT or a SAML assertion for an OAuth access token.

- 8. Define one or more OpenID Connect policies using the Applications # OAuth # OpenID Connect Policy Management window if you support OpenID Connect use cases.
- 9. Go to Applications # OAuth # Clients and create one or more OAuth clients in the Client window.
- 10. Optional: Configure client settings and registration policies for dynamic client registration.
- 11. Optional: Configure client session management settings.

Configuring authorization server settings

The Authorization Server Settings window provides controls over the usage and behavior of PingFederate as an OAuth authorization server, including the policies and settings for various grant types, refresh-tokens, persistent grants, and ID tokens.

Steps

- 1. Go to System # OAuth Settings # Authorization Server Settings.
- 2. Configure authorization server settings to suit your use cases.

The following table describes each field.

Field	Description
Authorization Code Timeout (Seconds)	The amount of time in seconds that an authorization code is considered valid. The default value is 60.
Authorization Code Entropy (Bytes)	The length in bytes of the authorization code returned to clients. The default value is 30.
Track User Sessions for Logout	When selected, PingFederate links the sessions for identity provider (IdP) adapters that are used by clients to the PingFederate authentication session of the resource owner. When a user initiates logout, PingFederate sends logout requests to close the adapter sessions through the browser.
	Selecting this check box also allows per-client logout endpoints to be defined, which will be invoked during logout.
	This check box is not selected by default.
Refresh Token and Persistent Grant Settings	

Field [Persistent Grant Max]

Lifetime

Description

This field determines whether persistent grants should expire, and if so, the default maximum lifetime for persistent grants.

- Select Grants Do Not Expire to let persistent grants remain valid until they are revoked or removed.
- Select the time value option to set a maximum lifetime for persistent grants. Enter an integer between 1 and 999 in the field and select Days, Hours, or Minutes from the list.

The default selection is **Grants Do Not Expire**.



Note:

You can override the expiration in individual client records or grant-mapping configurations. Grant-mapping configurations take precedence and require an extended persistent grant attribute, **PERSISTENT_GRANT_LIFETIME**.

Persistent Grant Idle Timeout

This field determines whether persistent grants should expire due to inactivity, regardless of whether the maximum lifetime has been reached, and the default idle timeout period.

- Select Grants Do Not Timeout Due To Inactivity and let persistent grants remain valid until they expire or are removed.
- Select the time value option to set the default idle timeout window.
 Enter an integer between 1 and 999 in the field and select a unit of measurement from the list.

If you configure an idle timeout value, the idle timeout window slides when a persistent grant updates. When you have an idle timeout value configured without a maximum lifetime, persistent grants remain valid until they expire due to inactivity or until the grant storage revokes or removes them. When you have an idle timeout value configured with a maximum lifetime, persistent grants remain valid until they expire due to inactivity or lifetime expiration or until the grant storage removes them. For more information, see *Transient grants and persistent grants* on page 79.

For new installations, the default inactivity allowance is 30 days. For upgrades, **Grants Do Not Timeout Due To Inactivity** is the default selection unless a specific value was set previously.



Note:

You can override the expiration in individual client records.

Refresh Token Length (Characters)

Refresh Token Length The length of the refresh tokens in numerical characters.

The default value is 42.

Field	Description
Roll Refresh Token Values (Default Policy)	When selected, PingFederate generates a new refresh token when a new access token is issued. Otherwise, each refresh token is used until it becomes invalid, either by manual revocation or another security setting that renders the token invalid.
	Note:
	New refresh tokens are not issued during the interval defined by the Minimum Interval to Roll Refresh Tokens field.
	This check box is not selected by default.
Minimum Interval to Roll Refresh Tokens (Hours)	The minimum number of hours that must pass before a new refresh token can be issued. This setting provides a way to allow for rolling refresh tokens without having to send a new refresh token on every request.
	The default value is 0.
Reuse Existing Persistent Access Grants for Grant Types	If a client makes multiple requests for the same user and the same or lesser scope, select the grant types that you want PingFederate to reuse the existing grant for, rather than creating a new grant for each request.
	Reusing an existing persistent grant imposes a limit of one grant per client, per user. In the context of refresh tokens, only the most recently issued is valid, and the previously issued refresh token is invalidated. If the same client are installed on multiple devices and used regularly by the same user, the grant type used by this client should be cleared.
	The applicable grant types are:
	 Implicit (default) Authorization Code Resource Owner Password Credentials
	When the Implicit check box is selected, PingFederate requests consent from the user only once. The user is not asked for authorization on subsequent requests until the access grant is revoked.
	When the Authorization Code check box is selected, the same is true if the Bypass Authorization for Previously Approved Persistent Grants check box is also selected.
Allow Unidentified Clients to Make	When selected, PingFederate allows resource owners to obtain access tokens without client ID or client authentication.
Resource Owner Password Credentials	The check box is not selected by default.

Grants

Field	Description
Allow Unidentified Clients to Request Extension Grants	When selected, PingFederate allows user-initiated or client-initiated events, such as a mobile application or a scheduled task, to obtain access tokens without the client presenting a client_id or client_secret for the extension grant types, namely:
	 JWT Bearer Token grant type
	urn:ietf:params:oauth:grant-type:jwt-bearerSAML 2.0 Bearer Assertion grant type
	<pre>urn:ietf:params:oauth:grant-type:saml2-bearer • Validation grant type</pre>
	urn:pingidentity.com:oauth2:grant_type:validate_beare
	The check box is not selected by default.
Token Endpoint Base URL	When clients authenticate with the private_key_jwt authentication method, PingFederate validates the aud parameter value found inside the signed JWT against its base URL or any configured virtual host names. If the values do not match, the authentication fails.
	Enter a separate base URL that PingFederate can take into consideration as well when validating the aud parameter.
	If configured, the OpenID Provider configuration endpoint /.well-known/openid-configuration uses the Token Endpoint Base URL field value as the base for the token endpoint.
	This field has no default value.
Persistent Grant Exten	ded Attributes

Field	Description
Consent User Interface	Specifies whether PingFederate or a trusted web application should handle consent approval.
	Default
	Select Default and let PingFederate handle consent approval by presenting the Request for Approval page to the resource owner.
	External
	Select External to delegate the responsibilities of consent approval to a trusted web application.
	For example, if you have created an instance of the Reference ID Adapter version (version 1.5 or a subsequent version), you can select it from the list. The expectation is that the trusted web application is integrated with PingFederate through this Reference ID Adapter instance.
	When selected, you must also configure two additional fields: External Consent IdP Adapter and External Consent Scopes Attribute.
	External Consent IdP Adapter
	The External Consent IdP Adapter field displays a list of IdP adapter instances that are capable of facilitating the consent approval process.
	For example, if you have created an instance of the Reference ID Adapter version (version 1.5 or a subsequent version), you can select it from the list. The expectation is that the trusted web application is integrated with PingFederate through this Reference ID Adapter instance.
	Your development team can also create a custom adapter using the PingFederate SDK. For more information, see the Javadoc for the IdpAuthenticationAdapterV2 interface, the ExternalConsentPageAdapter.java file for a sample implementation, and the SDK developer's guide for build and deployment information.
	Tip:
	The Javadoc for PingFederate and the sample implementation are in the <pf_install>/pingfederate/sdk directory.</pf_install>
	Once deployed, you can create an instance of the custom adapter and select it from the list.
	External Consent Scopes Attribute
	The Fixtage of Consent Consen Attribute field displays a list

The **External Consent Scopes Attribute** field displays a list of attributes defined in the IdP adapter contract of the selected adapter instance. Select the attribute whose value contains the approved scopes returned by the trusted web application.

For example, if you have added an attribute called **approvedScopes** to the adapter instance with the expectation that **approvedScopes** is the attribute that the trusted web application passes approved scopes to PingFederate, select **approvedScopes** from the list.

Field	Description		
OAuth Administrative	OAuth Administrative Web Services Settings		
Password Credential Validator	Selects a Password Credential Validator (PCV) instance to manage clients using the OAuth Client Management Service or to manage persistent grants using the OAuth Access Grant Management Service.		
	This setting has no default selection. When no PCV is selected, neither service can be used.		
Persistent Grant Mana	gement API		
management if the users approve a specific scope the client requests an ac token remains valid, the	anagement API allows clients to assume the responsibility of grant is authorize the clients to do so. In this scenario, a client prompts the user to be for managing persistent grants on the user's behalf. If the user approves, access token with such scope from PingFederate. As long as the access client can retrieve and revoke persistent grants and their associated as and values for that user. For more information, see OAuth Persistent Grant ge 1117.		
Access Token Manager	Selects an Access Token Management instance under which one or more clients can use the access tokens issued to manage persistent grants on their users' behalves.		
	Such clients must also be configured to use this Access Token Management instance in their client configurations.		
Required Scope	Selects the scope that PingFederate looks for in access tokens prior to granting clients the permission to manage persistent grants on their users' behalves.		
	Clients must obtain access tokens with this scope and include them in their grant-management requests.		

Cross-Origin Resource Sharing Settings

Field Description **Allowed Origin**

Enter any number of trusted origins to enable cross-origin resource sharing (CORS) support for the following OAuth endpoints:

- /as/token.oauth2
- /as/revoke token.oauth2
- /idp/userinfo.openid
- /pf-ws/rest/oauth/grants/
- /pf/JWKS
- /.well-known/openid-configuration
- /as/bc-auth.ciba

Once configured, client-side web applications from the trusted origins are allowed to make requests to the PingFederate authorization server for the purpose of accessing protected resources, such as obtaining or renewing access tokens with refresh tokens, presenting access tokens for revocation, querying additional claims (user attributes), and retrieving OpenID Provider configuration information and JSON Web Key Sets. For more information about CORS, see W3C's recommendation on Cross-Origin Resource Sharing.

Sample entry **Expected behavior** CORS requests originating from https:// https:// www.example.com are allowed. https:// CORS requests originating from https:// www.example.com:8080 are allowed. CORS requests originating from https:// https:// www.example.com:<any port> are allowed. Note: Given this sample entry, a port number is required in the Origin request header.



Important:

While using the wildcard character provides the convenience of allowing multiple origins with one entry, consider adding individual origins to limit CORS requests to a list of trusted hosts.

This field has no default entry. PingFederate allows multiple entries.

Device Authorization Grant Settings

The OAuth 2.0 Device Authorization Grant specification defines the process of a user granting authorization to a device using a browser on a second device, such as a smartphone or a computer. For more information, see *Device authorization grant* on page 51.

You can override this setting in individual client records.

Field Description Registered This field controls whether PingFederate should replace the user **Authorization Path** authorization endpoint with a different path, perhaps for ease of use or branding purposes, when formulating the verification URLs to be (Optional) included in its device authorization responses. For more information, see *Device authorization endpoint* on page 1077. The domain portion remains to be the base URL of PingFederate. For example, if the base URL is https://www.example.com and this field is configured with a value of /go, PingFederate returns https://www.example.com/go and https://www.example.com/go?user_code=<activationcode> as the user authorization URLs. If PingFederate receives a device authorization request at one of the configured virtual host names, PingFederate preserves the virtual host name in its device authorization responses. Note: The registered authorization path behaves the same way as the user authorization endpoint does. For more information, see *User authorization* endpoint on page 1080. This field is ignored when the User Authorization URL field is configured here or in the individual client records. The configured value must begin with a forward slash. This field has no default value. Pending Authorization The lifetime of an activation code (the user code parameter value) in Timeout (seconds) seconds. The default value is 600. You can override this setting in individual client records. Device Polling Interval The amount of time in seconds that the device waits between polling (seconds) requests to the PingFederate token endpoint. The default value is 5.

You can override this setting in individual client records.

Note:

External consent user interface

As use cases evolve and give users more control over their data, it is important to provide detailed information about the requests. In addition to scope descriptions, PingFederate supports the use of an external web application to prompt for authorization consent.

An external web application provides the opportunity to retrieve additional information specific to the users. For example, if a client requests the read_bank_account scope, the web application can retrieve the user's customer information file and give the user the ability to choose which accounts to be made available to the client.

To use an external web application for consent approval, configure the **Consent User Interface** setting in System # OAuth Settings # Authorization Server Settings. Select the External option and then configure the External Consent IdP Adapter and External Consent Scopes Attribute settings accordingly.

Responsibilities of the external web application

Delegating consent approval to an external web application implies that PingFederate can trust the web application. PingFederate expects this trusted web application to fulfill the following responsibilities:

- Retrieve from PingFederate the list of requested scopes in a secure manner.
 - For example, when integrating the web application with PingFederate through an instance of the Reference ID Adapter, such communications occur through a direct connection between the web application and PingFederate. This back-channel connection is protected by authentication and encryption (HTTPS).
- Provide to the resource owner the information associated with the list of requested scopes and the user interface elements to approve or deny the requested scopes.
- Validate that the approved scopes found in the response from the resource owner do not exceed the requested scopes.



Important:

This validation guards against unauthorized access in the event that the response is tampered and the original approved scopes are compromised.

- As needed, modify the approved scopes before returning them to PingFederate.
 - This allows the web application to override authorization decisions.
- Return the list of approved scopes to PingFederate in a secure manner.

By default, PingFederate handles consent approval by presenting the Request for Approval page to the resource owner. Upon receipt of the response from the resource owner, PingFederate validates that the approved scopes do not exceed the requested scopes. If the validation passes, PingFederate adds the approved scopes to the access token. Otherwise, PingFederate returns the invalid scope error to the client.

When an external consent user interface is enabled, PingFederate delegates consent approval to an external web application. As PingFederate trusts this web application, it always adds the scopes returned by the trusted web application to the access token, regardless of whether the returned scopes have already been defined in the system. The issuance of the access token is still subject to the criteria defined in the grant mapping configuration, the token mapping configuration, or both. For more information, see Grant contract mapping on page 528 and Token mapping on page 549.

Scopes and scope management

OAuth allows you to constrain the privileges associated with an access token, and scopes allow you to define the privileges requested and granted.

Static scopes versus dynamic scopes

As an authorization server, PingFederate supports the concepts of static scopes and dynamic scopes. To define a static scope, use a text value such as read bank account. To define a dynamic scope, use a text value with a variable component represented by a wildcard, such as read bank account txn:*. As illustrated, dynamic scopes allow clients to request authorization using scope values with a variable component from one request to another. For example, when a client sends an authorization or token request with a requested scope of read bank account txn:1234, PingFederate can match the requested scope to the dynamic scope pattern of read bank account txn: * and can issue an access token with the requested scope of read bank account txn:1234.

Scope groups

For ease of management and subsequent client interactions, PingFederate has the capability to create multiple groups of static scopes. A client can reference a scope group in applicable OAuth 2.0 protocol interactions. When authorized, clients can subsequently request access tokens with fewer permissions by presenting to the token endpoint a refresh token and the desired subset of scopes.

A scope group must contain at least one static scope, and multiple sub scopes are allowed. Multiple scope groups can share the same set of sub scopes. However, no scope group can contain another scope group or the default scope.

Scope group expansion

An authorization request can include one or more scope values. If the request is authorized, PingFederate issues an access token to the client. When the client brings the access token to a resource server to access protected resources, the resource server may contact PingFederate to validate the access tokens. Scope groups are not expanded in JSON web token (JWT)-based access tokens or token introspection responses by default. You can optionally enable scope group expansion per access token management instance.



Note:

Regardless of whether you choose to expand scope groups, the Request for Approval window always presents the description of the requested scope groups, if any.

PingFederate has the flexibility to manage common and exclusive scopes and scope groups.

Common scopes and scope groups

Common scopes and scope groups are optional. If defined, they are available to all clients by default. As needed, you can restrict individual clients to a subset of common scopes or scope groups in their configurations.

Clients created via the Dynamic Client Registration protocol can also be restricted to a subset of common scopes or scope groups based on the configuration on the Scope Constraints tab in System # OAuth Settings # Client Settings. The Scope Constraints configuration is shared across all clients registered through dynamic client registration. If a certain client requires a different set of common scopes or scope groups, modify the client configuration by using the administrative console, the administrative API, or the OAuth Client Management Service after the client has been created.

Exclusive scopes and scope groups

Exclusive scopes and scope groups are optional. If defined, they are restricted from all clients by default. As needed, you can configure individual clients to allow a subset of exclusive scopes or scope groups in their configurations.

Clients created with the Dynamic Client Registration protocol can also be configured to allow a subset of exclusive scopes or scope groups based on the Scope Constraints tab on the System # OAuth Settings # Client Settings window. The Scope Constraints configuration is shared across all clients registered via dynamic client registration. If a certain client requires a different set of exclusive scopes or scope groups, modify the client configuration by using the administrative console, the administrative API, or the OAuth Client Management Service after the client has been created.



A scope or scope group is either a common scope or group, or an exclusive scope or group. Duplicate scopes and scope groups are not allowed. Scope and scope group values are case-sensitive.



Create scopes that are intended for the majority of clients as common scopes. Create scopes that should be limited to the minority of clients as exclusive scopes.

OpenID Connect

If one or more clients support the OpenID Connect standard, add the following scopes for the purpose of requesting specific sets of claims from the OpenID Provider:

- openid
- address
- email
- phone
- profile



Z Tip:

If most clients are allowed to use these scopes, create them as common scopes.

Per-client scope management

You can manage scope access on a client-to-client basis. The client settings are **Restrict Common** Scopes and Exclusive Scopes.

Restrict Common Scopes

This setting determines whether all common scopes and scope groups should be made available to the client, or only a select few.

When selected, the administrative console displays a list of existing common scopes and scope groups. Choose the common scopes and scope groups that are intended for the client. The rest, as well as any future common scopes and scope groups, become invalid for the client. If the client tries to use such scope or scope group, it will receive an invalid scope error message from PingFederate.

When cleared, all existing and future common scopes and scope groups are available to the client. This is the default behavior.

Exclusive Scopes

This setting determines whether any exclusive scopes and exclusive scope groups should be made available to the client.

When selected, the administrative console displays a list of existing exclusive scopes and scope groups. Choose the exclusive scopes and scope groups that are intended for the client. The rest, as well as any future exclusive scopes and scope groups, become invalid for the client. If the client tries to use such scope or scope group, it will receive an invalid scope error message from PingFederate.

When cleared, no exclusive scopes and scope groups are available to the client. This is the default behavior.



Both settings impact dynamic scope evaluation. For more information, see *Dynamic scope evaluation and* per-client scope management.

Dynamic scopes

A dynamic scope is defined by using a text value with a variable component represented by an asterisk (*). PingFederate supports three dynamic scope patterns:

- A prefix followed by a wildcard, for example: prefixTextValue*
- A wildcard followed by a suffix, for example: *suffixTextValue
- A wildcard placed between a prefix and a suffix, for example: prefixTextValue*suffixTextValue

PingFederate only allows one variable component. Backslashes (\) and double quotation marks (") are not allowed in the prefix or the suffix. Multiple dynamic scopes are supported in conjunction with any number of static scopes and scope groups.

Dynamic scope evaluation

When a client sends an authorization or token request with a list of desired scopes, PingFederate validates the requested scopes against its configurations.

If PingFederate finds no match for the requested scopes, it returns an invalid scope error message to the client.

If PingFederate matches the requested scope to an existing static scope or scope group, it checks the client configuration to determine whether such static scope or scope group is valid for the client. If it is, PingFederate proceeds further. For example, if PingFederate is configured to handle consent approval, it presents to the user the Request for Approval window with the description associated with the matched static scope or scope group. If PingFederate should issue an access token, the token is issued with the requested scope. If such static scope or scope group is not valid for the client, PingFederate returns an invalid scope error message to the client.

If PingFederate finds no exact match but finds a partial match to one or more dynamic scopes, the partial match with the highest number of matched characters in the prefix, suffix, or both is the matched dynamic scope. In the event that two partial matches tie, the partial match with the highest number of characters matched in the prefix is the matched dynamic scope. PingFederate then checks the client configuration to determine whether such dynamic scope is valid for the client. If it is, PingFederate proceeds further. Otherwise, PingFederate returns an invalid scope error message to the client. If PingFederate should issue an access token, the token is issued with the requested scope, not the matched dynamic scope pattern.

Example

For example, you add the following dynamic scopes:

Common Scopes	Exclusive Scopes
*123	zSomeExclusiveScope
*12345	
a*c#123	
ab*#123	
xy*123	
xy*	

You also add a client without any common scope restrictions. This client can access all common scopes.

The following table illustrates the expected results when the client sends an authorization request with these scopes:

- xy#1
- xy#12
- xy#123
- xy#1234
- xy#12345
- xy#123456
- xyz
- z123
- z12345
- abc#123

Requested scope	Matched dynamic scope	Variable component from the requested scope
xy#1	xy*	#1
xy#12	xy*	#12
xy#123	xy*123	#
xy#1234	xy*	#1234
xy#12345	*12345	xy#
xy#123456	xy*	#123456

The minimum length of the variable component is one character. If the variable component contains two or more characters, it may also contain the asterisk character as well. Given the same common dynamic scopes and the same client configuration, requested scopes of xyQ123, xy*Q123, xy*Q123, xy*YQ123, xy*YQ123 will be matched as in the following table.

Requested scope	Matched dynamic scope	Variable component from the requested scope
xyQ123	xy*123	Q
xy*Q123	xy*123	*Q
xyQ*123	xy*123	Q*
xy**Q*123	xy*123	**Q*

If the client sends an authorization request with a requested scope of xy*123, it will receive an invalid scope error from PingFederate.

Dynamic scope evaluation and per-client scope management

Depending on the configured dynamic scope patterns and if they are defined as common or exclusive dynamic scopes, per-client scope management settings can impact the results of scope evaluation.

The **Restrict Common Scopes** setting determines whether all common scopes and scope groups are available to the client, or only a select few. Use this setting to restrict certain common dynamic scopes.

The **Exclusive Scopes** setting determines whether any exclusive scopes and scope groups are available to the client. When this check box is not selected, PingFederate does not consider any exclusive dynamic scopes or any exclusive static scopes and scope groups when trying to match a requested scope against a list of configured scopes and scope groups. When the check box is selected, all exclusive scopes and scope groups are considered. If PingFederate matches a requested scope to an exclusive dynamic scope and such scope is not available to the client, PingFederate returns an <code>invalid_scope</code> error message to the client. This remains true for a lesser partial match to an available common dynamic scope.

Example

For example, you update your previous sample scope configuration as in the following table.

Common Scopes	Exclusive Scopes
*123	xy*123
*12345	zSomeExclusiveScope
a*c#123	
ab*#123	
xy*	

The following table describes the results when the client sends an authorization request with a requested scope of xy#123.

Per-client scope management settings	Result	Variable component from the requested scope
Restrict Common Scopes	PingFederate matches the requested	se
Not selected.	scope of xy#123 to the common dynamic scope of *123.	
Exclusive Scopes	The exclusive dynamic scope of xy*123	
Not selected.	is not taken into consideration because the Exclusive Scopes check box is not selected. PingFederate does not consider any exclusive scopes and scope groups as eligible candidates in this scenario.	
Restrict Common Scopes	PingFederate returns an	Not applicable.
 Not selected. 	invalid_scope error message, because the exclusive dynamic scope	
Exclusive Scopes	of xy*123 is not allowed based on the	
 Selected. 	Exclusive Scopes configuration.	
 zSomeExclusiveScope is selected. 		
No other exclusive scope is selected.		
Restrict Common Scopes	PingFederate matches the requested	#
Not selected.	scope of xy#123 to the exclusive dynamic scope of xy*123.	
Exclusive Scopes	5 ,	
 Selected. 		
xy*123 is selected.		
It does not matter if any other exclusive scope is selected.		
Restrict Common Scopes PingFederate returns an		Not applicable.
Selected.	invalid_scope error message, because the common dynamic	
 xy* is selected 	scope of *123 is not allowed based	
No other common scope is selected.	on the Restrict Common Scopes configuration.	
Exclusive Scopes	The exclusive dynamic scope of xy*123 is not taken into consideration because	
Not selected.	the Exclusive Scopes check box is not selected. PingFederate does not consider any exclusive scopes and scope groups as eligible candidates in this scenario.	

Per-client scope management settings	Result	Variable component from the requested scope
Restrict Common Scopes	PingFederate matches the requested	#
Selected.	scope of xy#123 to the exclusive dynamic scope of xy*123, because	
*123 is selected	xy*123 is a better partial match than	
It does not matter if any other common scope is selected.	*123.	
Exclusive Scopes		
Selected.		
 xy*123 is selected. 		
It does not matter if any other exclusive scope is selected.		

Description for scopes and scope groups

When defining a scope or a scope group, enter a value and a description for the scope or the scope group. This description helps you identify the purpose of the scope or scope group at a later time. If PingFederate is configured to handle consent approval, the **Scope Description**, **Scope Group Description**, and **Default Scope Description** fields determines the text that appears on the **Request for Approval** window.

Default scope

The default scope is the implied permissions when no scope and scope group values are indicated, or in addition to any scope or scope group values.

If your organization requires a localized description, enter a unique alias in the **Default Scope Description** field, such as oauth.approval.page.template.defaultScope. Insert the same alias with the desired localized text in the applicable language resource files, located in cpf install/pingfederate/server/default/conf/language-packs.

Static scopes and scope groups

You can enter simple descriptions or localize the descriptions by using the PingFederate localization framework.

Dynamic scopes

You can enter simple descriptions. You can also use a mix of text and scope-description variables. \${scope} represents the requested scope, and \${scope-var} represents the variable component found in the requested scope.

Suppose you added a dynamic scope with a pattern of dynaGet67*10 and a scope description of \${scope} contains \${scope-var}. If a client requests a scope value of dynaGet67eight910, the resulting scope description is dynaGet67eight910 contains eight9. (eight9 is the variable component found in the requested scope.)

If your organization requires a localized description, enter a unique alias in the **Scope Description** field (for example, oauth.approval.page.template.someDynamicScope). Then insert the same alias with the desired localized text in the applicable language resource files, located in the $< pf_install > /pingfederate/server/default/conf/language-packs directory. You may also use scope-description variables as part of the localized text.$



Dynamic scopes and consent user interface

The default consent approval process and user interface in PingFederate are capable of handling dynamic scopes and their scope descriptions. While the scope description and the optional scope-description variables provide the basic controls to describe a given scope, PingFederate also supports the use of an external web application to prompt for authorization consent. This allows you to retrieve additional information specific to the users and apply application-specific scope-processing logic.

Coordinating with developers

Regardless of whether a static scope, a scope group, or a dynamic scope is created as common or exclusive, a scope or a scope group represents access to a resource or API on the RS. Applicable scope or scope group values require coordination with developers that are familiar with the details of the RS OAuth implementation. For clients supporting the OpenID Connect protocol, you can direct the developers to your PingFederate OpenID Provider configuration endpoint to retrieve a list of common scopes and common scope groups.



The OpenID Provider configuration endpoint does not return exclusive static scopes, exclusive scope groups, common dynamic scopes, and exclusive dynamic scopes by default. You can optionally customize the response to include such scopes and scope groups as needed.

Defining scopes

PingFederate manages scopes and scope groups in common and exclusive buckets.

About this task

Common scopes and scope groups are optional. They are available to all clients by default. You can restrict individual clients to a subset of common scopes or scope groups on a client-by-client basis in their client configurations.

Exclusive scopes and scope groups are also optional. They are restricted from all clients by default. However, you can grant individual clients access to one or more exclusive scopes or scope groups in their client configurations.

You can also create static scopes, static scope groups, and dynamic scopes. Scope groups allow clients to request a "super scope" and optionally downgrade to a subset of it later. Dynamic scopes address the business requirement where clients want to request authorization by using scope values with a variable component from one request to another. For detailed information about scopes, see Scopes and scope management on page 486.

You can manage scopes, scope groups, and the default scope description from System # OAuth Settings on the Scope Management window. Configuration steps for common and exclusive scopes and scope groups are identical.



Note:

A scope or scope group is either a common scope or group, or an exclusive scope or group. Duplicate scopes and scope groups are not allowed.

Tip:

Create scopes that are for the majority of clients as common scopes. Create scopes that are for a minority of clients as exclusive scopes. You can organize common or exclusive static scopes into common or exclusive scope groups.

Steps

- 1. Go to System # OAuth Settings # Scope Management.
- 2. On the Common Scopes or Exclusive Scopes tab, configure any number of common or exclusive static scopes.
 - a. Enter a static scope value and a description in the **Scope Value** and **Scope Description** fields.

Scope values are case-sensitive. A requested scope value of email does not match a configured static scope value of Email. Do not use backslashes (\) or double quotation marks (") in the Scope Value field.



Note:

If PingFederate is configured to handle consent approval and your organization requires a localized description, enter a unique alias in the **Scope Description** field. For more information, see Description for scopes and scope groups.

- b. Click Add.
- c. Repeat to configure additional static scopes. Display order does not matter.
- 3. On the Common Scopes or Exclusive Scopes tab, configure any number of common or exclusive dynamic scopes.
 - a. Enter a dynamic scope pattern and a description in the Scope Value and Scope Description fields.

You must use a case-sensitive text value with the asterisk character (*). For supported dynamic scope patterns, see Dynamic Scopes.

You can enter a simple description, or you can use a mix of text and scope-description variables. For more information, see *Description for scopes and scope groups*.



Note:

If your organization requires additional information specific to the users or application-specific scope-processing logic, you can configure PingFederate to use an external web application to handle consent approval.

- b. Select the check box under **Dynamic**.
- c. Click Add.
- Repeat to define additional dynamic scopes. Display order does not matter.

- 4. On the Common Scopes or Exclusive Scopes tab, configure any number of common or exclusive scope groups.
 - a. Enter a scope group value and its description under Scope Group Value and Scope Group Description.

If PingFederate is configured to handle consent approval and your organization requires a localized description, enter a unique alias in the Scope Group Description field. For more information, see Description for scopes and scope groups.



Note:

All scope groups are defined as static scope groups. The partial-matching concept is intended for dynamic scopes only and does not apply to scope groups. Like static scope values, scope group values are case-sensitive. Do not use backslashes (\) or double quotation marks (") in the **Scope** Group Value field.

Select at least one static scope under Sub Scopes.



Note:

The administrative console filters out dynamic scopes because the scope-grouping capability is reserved for static scopes only.

- c. Click Add.
- d. Repeat to define additional scope groups. Display order does not matter.
- 5. Optional: On the Common Scopes or Exclusive Scopes tab, use the Edit, Update, and Cancel buttons to make or undo a change to an existing entry. Use the Delete and Undelete buttons to remove an existing entry or cancel the removal request.



CAUTION:

Updating or removing a scope or scope group value in production can cause runtime errors when a client request specifies the scope or scope group using the previously defined value. Errors can also occur when the requesting client is restricted to a scope or scope group that no longer exists unless the affected client configuration is also updated.

6. On the **Default Scope** tab, enter a description for the default scope.

If PingFederate is configured to handle consent approval and your organization requires a localized description, enter a unique alias in the **Default Scope Description** field. For more information, see Description for scopes and scope groups.



Note:

The default scope is the implied permissions when no scope and scope group values are indicated, or in addition to any scope or scope group values.

7. Click Save.

Result

Scopes and scope groups represent access to resources or APIs on the resource server . For clients supporting the OpenID Connect protocol, you can direct the developers to your PingFederate OpenID Provider configuration endpoint to retrieve a list of common scopes and common scope groups.

Note:

By default, the OpenID Provider configuration endpoint does not return exclusive static scopes, exclusive scope groups, common dynamic scopes, and exclusive dynamic scopes. You can customize the response to include such individual scopes and scope groups.

Configuring client settings

Use the Client Settings window to configure dynamic client registration settings.

Steps

- 1. On the **Dynamic Client Registration** tab, modify settings as needed.
- 2. Click **Next** to advance to the next tab.
- 3. Click Save to retain your changes.

Configuring dynamic client registration settings

Dynamic client registration allows developers to register OAuth clients through an API based on open standards.

About this task

PingFederate supports various client metadata as described in Supported client metadata on page 498. If specific use cases require additional metadata, add them as extended properties in System # Server # **Extended Properties.**



Important:

As dynamic client registration can expose your server to unwanted client registrations, we recommend protecting PingFederate by requiring an initial access token, configuring one or more client registration policies, and protecting access to the dynamic client registration endpoint.



Note:

Dynamic client registration requires OAuth client storage in an external datastore, such as a database or LDAP directory. If you have not yet switched from the default on-disk client storage to an external datastore, see *OAuth client datastores* on page 895. You can continue with the rest of the configuration; however, dynamic client registration remains inactive until an external client storage is defined.

Steps

- 1. Go to System # OAuth Settings # Client Settings and click Dynamic Client Registration.
- 2. Select the check boxes for the options that you want to use.

The following table describes each option.

Option	Description
Enable Dynamic Client Registration	Select this option to enable dynamic client registration. This option is disabled by default.

Option	Description
Require Initial Access Token	Select this option to require an initial access token. If selected, you must also select the required scope or scope group from the list.
	Important:
	Although optional, selecting this option can prevent unwanted client registrations.
	Note:
	Developers must be set up to obtain access tokens with the required scope or scope group from your PingFederate authorization server. For example, you can create a new OAuth client for a group of developers, assign this client a specific scope for the purpose of creating other clients using the OAuth 2.0 Dynamic Client Registration protocol, and let the developers obtain their access tokens directly by completing one of the supported OAuth flows. You can also write a custom web application that uses the OAuth flow to obtain access tokens on behalf of the developers as they make their requests.
	This option is enabled by default.
Enable Dynamic Client Registration Management	Select this option to enable dynamic client registration management, and to make the following three client management/maintenance options visible.
	 Rotate Client Secret - Select this option to rotate the client secret when a client updates or retrieves its configuration. This option is enabled by default. Rotate Registration Access Token - Select this option to rotate the registration access token when a client updates or retrieves its configuration. This option is enabled by default. Allow Client Delete - Select this option to allow clients to deprovision themselves on the authorizations server. This option is enabled by default.
	Dynamic client registration management allows a client to retrieve its configuration through GET requests, to update its configuration through PUT requests to the provided registration client URI, and to deprovision itself through a DELETE.
	For more information, see the <i>OAuth 2.0 Dynamic Client Registration Management Protocol</i> .
	This option is disabled by default.

Result

When dynamic client registration is active, developers can send client registrations to the /as/clients.oauth2 endpoint to create OAuth clients dynamically.

Other maintenance calls can be made to the $registration_client_uri$ returned in the original registration response. This endpoint has the format /as/clients.oauth2/<clientId>.

Supported client metadataPingFederate supports various client metadata, as described in the following table.

Metadata field	Metadata description
client_name	A descriptive name for the client instance. This name appears when the user is prompted for authorization.
token_endpoint_auth	_methione client authentication method.
	PingFederate accepts the following values:
	noneclient_secret_basicclient_secret_posttls_client_auth
	For more information, see <i>Mutual TLS Profiles for OAuth clients</i> . • private_key_jwt
	For more information, see Client Authentication.
tls_client_auth_subje	ct_dnThe subject DN of the client certificate.
	This field is required if tls_client_auth is the value of the token_endpoint_auth_method parameter.
token_endpoint_auth	_signi Fige alig ning algorithm that the client must use to sign the JSON web tokens (JWT) for client authentication.
	This field applies only when the token_endpoint_auth_method parameter is provided with a value of private_key_jwt.
	PingFederate accepts the following values:
	 RS256 - RSA using SHA-256 RS384 - RSA using SHA-384 RS512 - RSA using SHA-512 ES256 - ECDSA using P256 Curve and SHA-256 ES384 - ECDSA using P384 Curve and SHA-384 ES512 - ECDSA using P521 Curve and SHA-512 PS256 - RSASSA-PSS using SHA-256 PS384 - RSASSA-PSS using SHA-384 PS512 - RSASSA-PSS using SHA-512
	Note:
	RSASSA-PSS signing algorithms require either a Java 11 runtime environment or an integration with a hardware security module (HSM) and a static-key configuration for OAuth and OpenID Connect. For more information on HSM integration and static keys, see and, respectively.
	If this parameter is not provided, the client can use any of the supported signing algorithms.

Metadata field

Metadata description

request object signing ald he signing algorithm that the client must use to sign its request objects for transmission of request parameters.

> Applicable only when the client might send its authorization requests using request objects.

PingFederate accepts the following values:

- RS256 RSA using SHA-256
- RS384 RSA using SHA-384
- RS512 RSA using SHA-512
- ES256 ECDSA using P256 Curve and SHA-256
- ES384 ECDSA using P384 Curve and SHA-384
- ES512 ECDSA using P521 Curve and SHA-512
- PS256 RSASSA-PSS using SHA-256
- PS384 RSASSA-PSS using SHA-384
- PS512 RSASSA-PSS using SHA-512



Note:

RSASSA-PSS signing algorithms require either a Java 11 runtime environment or an integration with a hardware security module (HSM) and a static-key configuration for OAuth and OpenID Connect. For more information on HSM integration and static keys, see and, respectively.

When this parameter is not provided, the client can use any of the supported signing algorithms.

For more information about request objects, see OpenID Connect specifications.

jwks_uri, and iwks

The URL of the JSON Web Key Set (JWKS) or the actual JWKS from the client.

If the client is configured to use the private_key_jwt client authentication method, to transmit request parameters in signed request objects, or to transmit CIBA request parameters in signed request objects, only one of the previous values is required for PingFederate to verify the authenticity of the JWTs.

Either value can be defined even if the client is not configured to use JWTs for authentication or transmission of request parameters. This flexibility allows the client to transmit request parameters in signed request objects for some requests and without the use of signed request objects for some other transactions. For information on runtime processing, see .

If the client signs its JWTs using an RSASSA-PSS signing algorithm, PingFederate must be deployed to run in a Java 11 runtime environment, or be integrated with a hardware security module (HSM) to process the digital signatures. For more information on HSM integration, see .

If the client is configured to encrypt ID tokens using an asymmetric encryption algorithm, either the JWKS URL or the actual JWKS must be provided. See the ID Token Key Management Encryption Algorithm setting.

Metadata field	Metadata description	
redirect_uris	An array of one or more redirection URIs where the OAuth AS may redirect the resource owner's user agent after authorization is obtained. The authorization code and implicit grant types require at least one redirection URI.	
logo_uri	The location of the logo used on user-facing OAuth grant authorization and revocation pages. For best results with the installed HTML templates, the recommended size is 72 x 72 pixels.	
scope	A space-separated list of one or more scopes, which a client can request.	
grant_types	An array of one or more grant types, which a client can request. PingFederate accepts the following values:	
	 authorization_code implicit refresh_token client_credentials urn:ietf:params:oauth:grant-type:device_code urn:openid:params:grant-type:ciba password extension (JWT Bearer Token or SAML 2.0 Bearer Assertion) 	
	For more information about each grant type, see Grant types.	

Metadata field	Metadata description	
response_types	An array of one or more response	e types, which a client can request.
	PingFederate accepts the following	ng values:
	 code code id_token code id_token token code token id_token id_token token token For more information about these Valued Response Type Combination and the specified response of the specified response types are to send one of the specified response types are the specified respon	e response types, see Definitions of Multiple- itions. e specified, the resulting client is only allowe onse types at runtime. Requests from this
	response type	grant types
	code	authorization_code
	code id_token	<pre>authorization_code and implicit</pre>
	godo id tokon tokon	
	code id_token token	<pre>authorization_code and implicit</pre>
	code token	-
	_	<pre>implicit authorization_code and</pre>
	code token	<pre>implicit authorization_code and implicit</pre>

Metadata field

Metadata description

id token signed responseTategJSON Web Signature (JWS) algorithm required for the OpenID Connect ID tokens.

Allowed values:

- none No signing algorithm
- HS256 HMAC using SHA-256
- HS384 HMAC using SHA-384
- HS512 HMAC using SHA-512
- ES256 ECDSA using P256 Curve and SHA-256
- ES384 ECDSA using P384 Curve and SHA-384
- ES512 ECDSA using P521 Curve and SHA-512
- RS256 RSA using SHA-256
- RS384 RSA using SHA-384
- RS512 RSA using SHA-512
- PS256 RSASSA-PSS using SHA-256
- PS384 RSASSA-PSS using SHA-384
- PS512 RSASSA-PSS using SHA-512



Note:

RSASSA-PSS signing algorithms require either a Java 11 runtime environment or an integration with a hardware security module (HSM) and a static-key configuration for OAuth and OpenID Connect. For more information on HSM integration and static keys, see and, respectively.



Important:

If static keys for OAuth and OpenID Connect are enabled, use either an RSA algorithm or an EC algorithm that has been configured with an active static key.

id_token_encrypted_respointse_allogorithm used to encrypt or otherwise determine the value of the content encryption key.

Allowed values:

- dir Direct Encryption with symmetric key
- A128KW AES-128 Key Wrap
- A192KW AES-192 Key Wrap
- A256KW AES-256 Key Wrap
- A128GCMKW AES-GCM-128 key encryption
- A192GCMKW AES-GCM-192 key encryption
- A256GCMKW AES-GCM-256 key encryption
- ECDH-ES ECDH-ES
- ECDH-ES+A128KW ECDH-ES with AES-128 Key Wrap
- ECDH-ES+A192KW ECDH-ES with AES-192 Key Wrap
- ECDH-ES+A256KW ECDH-ES with AES-256 Key Wrap
- RSA-OAEP RSAES-OAEP

Metadata field

Metadata description

id token encrypted responding compared encryption algorithm used to perform authenticated encryption on the plain text payload of the token.

> Required if an algorithm is provided through the id_token_encrypted_response_alg parameter.

Allowed values:

- A128CBC-HS256 Composite AES-CBC-128 HMAC-SHA-256
- A192CBC-HS384 Composite AES-CBC-192 HMAC-SHA-384
- A256CBC-HS512 Composite AES-CBC-256 HMAC-SHA-512
- AES-GCM-128 A128GCM
- AES-GCM-192 A192GCM
- AES-GCM-256 A256GCM

backchannel token deliverthertoklen delivery method that the client supports. PingFederate supports poll and ping.

> Set to poll if the client can check for the authorization results periodically at the token endpoint.

Set to ping if the client prefers to wait for a ping callback message from PingFederate as a signal that the authorization result is ready for pickup.

If this parameter is not provided and the CIBA grant type is enabled, the poll method is assumed.

backchannel_client_notificationclient(sometis) backchannel_client(sometis) backchannel back messages.

Required only if ping is the configured token delivery method.

backchannel_authenticationhesiqueingside in backchannel_authenticationhesiqueing in backchannel_authenticationhesiqueing in backchannel authenticationhesiqueing in backchannel authenticationhesi transmission of request parameters.

PingFederate accepts the following values:

- RS256 RSA using SHA-256
- RS384 RSA using SHA-384
- RS512 RSA using SHA-512
- ES256 ECDSA using P256 Curve and SHA-256
- ES384 ECDSA using P384 Curve and SHA-384
- ES512 ECDSA using P521 Curve and SHA-512
- PS256 RSASSA-PSS using SHA-256
- PS384 RSASSA-PSS using SHA-384
- PS512 RSASSA-PSS using SHA-512



Note:

RSASSA-PSS signing algorithms require either a Java 11 runtime environment or an integration with a hardware security module (HSM) and a static-key configuration for OAuth and OpenID Connect. For more information on HSM integration and static keys, see and, respectively.

If this parameter is not provided and the CIBA grant type is enabled, the client can use any of the allowed signing algorithms.

Metadata field	Metadata description	
backchannel_user_code_padees whether the client supports user code.		
	The purpose of this code is to authorize the transmission of an authentication request to the user's authentication device.	
	A valid value is either true or false.	
	If this parameter is not provided and the CIBA grant type is enabled, user code support is not enabled.	
	Note:	
	When user code support is enabled, the associated CIBA request policy must also be user code enabled.	
sector_identifier_uri	A URL using the HTTPS scheme that references a JSON file containing an array of redirect_uri values. For more information, see https://openid.net/specs/openid-connect-registration-1_0.html#SectorIdentifierValidation .	
subject_type	The type of subject used by the sector identifier, such as public or pairwise.	

Configuring scope constraints

On the **Scope Constraints** tab, you can configure which scopes or scope groups that developers can request when registering clients using dynamic client registration.

About this task

All clients created through dynamic client registration share this configuration. If a certain client requires a different set of common scopes, exclusive scopes, or both, modify the client configuration using the administrative console, the administrative API, or the OAuth Client Management Service after the client has been created. Scopes can also be overridden by client registration policies enforced during dynamic client registration.

Steps

- 1. Go to System # OAuth Settings # Client Settings and click Scope Constraints.
- To restrict clients created with the Dynamic Client Registration protocol to a subset of common scopes, select the **Restrict Common Scopes** check box and one or more applicable common scopes.

Result:

Your selections impact the developers in several ways:

- If you do not select the Restrict Common Scopes check box, developers can send client registrations without including the desired scopes. If the requests are valid, the clients are configured with all the common scopes and scope groups.
- If you select the Restrict Common Scopes check box without selecting at least one common scope or scope group, clients resulting from valid client registrations are configured without any common scopes or scope groups.
- If you select the Restrict Common Scopes check box with one or more applicable common scopes or scope groups, developers must send client registrations with the desired common scopes and scope groups. Otherwise, clients resulting from otherwise valid requests are also configured without any common scopes or scope groups.

 To allow clients created with the Dynamic Client Registration protocol to request for a subset of exclusive scopes, select one or more applicable exclusive scopes in the Allowed Exclusive Scopes field.

Result:

Your selections impact the developers in several ways:

- If you do not select any exclusive scope, clients resulting from valid client registrations are configured without any exclusive scopes or scope groups.
- If you select one or more applicable exclusive scopes or scope groups, developers must send
 client registrations with the desired exclusive scopes and scope groups. If they fail to do so, clients
 resulting from otherwise valid requests are also configured without any exclusive scopes or scope
 groups.

Result

Restricting common scopes and allowing exclusive scopes are not mutually exclusive. You can configure both options based on your use cases.

If you configure both options, developers must send client registrations with the desired common and exclusive scopes.

Depending on the configured dynamic scope patterns and whether they are defined as common or exclusive dynamic scopes, this configuration can impact the results of scope evaluation. The default scope is always available to all clients. For more information, see the **Dynamic scope evaluation and perclient scope management** section in *Scopes and scope management* on page 486.

Managing client configuration defaults

On the **Client Configuration Defaults** tab, specify the default settings that are proprietary to PingFederate for clients created with the OAuth 2.0 Dynamic Client Registration protocol.

About this task

Although these settings are shared among all clients created through dynamic client registration, they can be overridden by client registration policies enforced during dynamic client registration. You can also modify the client configuration using the administrative console, the administrative API, or the OAuth Client Management Service after the client has been created.

Steps

- 1. Go to System # OAuth Settings # Client Settings and click Client Configuration Defaults.
- 2. Optional: Modify the default values as needed.

The following table describes each field.

Field	Description
Private Key JWT - Replay Prevention	This field determines whether PingFederate mandates a unique signed JWT from the client for each request when the client is configured to authenticate via the private_key_jwt client authentication method, to transmit request parameters using in signed request objects, or to do both. This check box is not selected by default.
	Note: The underlying Assertion Replay Prevention Service is cluster-aware. For more information, see .

Field	Description
Require Signed Request	Determines whether the client must transmit request parameters in a single, self-contained parameter. The parameter name is request. The value of the request parameter is a signed JWT whose claims represent the request parameters of the authorization request. The OpenID Connect specification calls this JWT a request object.
	This check box is not selected by default.
Default Access Token Manager	Determines the default Access Token Management (ATM) instance for this client.
Persistent Grants Max Lifetime	Overrides the Persistent Grant Max Lifetime value set globally in System # OAuth Settings # Authorization Server Settings .
	Select one of the following options:
	 Use Global Setting (default) Grants Do Not Expire A custom value in days, hours, or minutes.
	Note:
	This setting can be overridden per grant-mapping configuration through the use of an extended persistent grant attribute
	PERSISTENT_GRANT_LIFETIME. The PERSISTENT_GRANT_LIFETIME attribute is defined in System # OAuth Settings # Authorization Server Settings . When this attribute is active, you can set the lifetime of persistent grants based on the outcome of attribute mapping expressions in individual grant-mapping configurations. For grant-mapping configurations that do not require this fine-grain control, you can configure them to use the default value.
Persistent Grants Idle	Overrides the Persistent Grant Idle Timeout field value set globally in System # OAuth Settings # Authorization Server Settings.
	Select one of the following options:
	 Use Global Setting (default) Grants Do Not Timeout Due To Inactivity A custom value in days, hours, or minutes.
	If you configure an idle timeout value, the idle timeout window slides when a persistent grant updates. When you have an idle timeout value configured without a maximum lifetime, persistent grants remain valid until they expire due to inactivity or until the grant storage revokes or removes them. When you have an idle timeout value configured with a maximum lifetime, persistent grants remain valid until they expire due to inactivity or lifetime expiration or until the grant storage removes them. For more information, see <i>Transient grants and persistent grants</i> on page 79.

Field	Description
Client Authentication Certificate Issuer DN	Select a trusted CA from the list. You can review CA certificates imported into PingFederate them in Security # Certificate & Key Management # Trusted CAs . You can select Trust Any to trust all the issuers found in the list.
	The default selection is None (Client TLS Certificate Authentication Disabled) , which does not allow developers to submit client registrations with a token_endpoint_auth_method parameter value of tls_client_auth.
Refresh Token Rolling Policy	Overrides the Roll Refresh Token Values setting configured globally in System # OAuth Settings # Authorization Server Settings.
	Select one of the following options:
	Use Global Setting (default)Roll
	This value does not override the Minimum Interval to Roll Refresh Tokens (Hours) value set in the Authorization Server Settings window. • Don't Roll

Policy

Select a specific OpenID Connect policy from the list.

Field Description **Device Authorization** Determines whether to use global device authorization grant settings defined in System # OAuth Settings # Authorization Server Settings. The default selection is Use Global Settings. You can select **Override** and configure any of the following settings. User Authorization URL This field determines whether PingFederate should use a different URL when formulating the verification URLs to be included in its device authorization responses. For more information, see . For example, if this field is configured with a value of https:// www.example.org/welcome, PingFederate returns https:// www.example.org/welcome and https://www.example.org/welcome? user_code=<activationcode> as the verification URIs. After processing the device authorization response, which includes the verification URIs, the device presents one of them to the user. The user is expected to browse to the presented verification URI on a second device. Important: The target web server must redirect the browser to PingFederate at its user authorization endpoint. For more information, see . It must also preserve the user code parameter value, if provided. For example, if your PingFederate server's base URL is https:// www.example.com and the User Authorization URL value is https://www.example.org/welcome, the target web server must redirect as follows: https://www.example.org/welcome to https://www.example.com/ as/user authz.oauth2 https://www.example.org/welcome?user_code=<activationcode> to https://www.example.com/as/user_authz.oauth2? user code=<activationcode> This field has no default value. Pending Authorization Timeout (seconds) The lifetime of an activation code (the user code parameter value) in seconds. This field has no default value. Device Polling Interval (seconds) The amount of time in seconds that the device waits between polling requests to the PingFederate token endpoint. This field has no default value. Bypass Activation Code Confirmation When PingFederate receives a verification request that includes an activation code (the user code parameter value), it prompts the user to confirm the activation code. For PingFederate to skip this confirmation step, select this check box. This check box is not selected by default.

Field	Description
Require Proof Key for Code Exchange (PKCE)	This field is only applicable when the client is configured to support the authorization code grant type.
	Determines whether the client must provide certain parameters to reduce the risk of authorization code interception attack. For more information, see the Proof Key for Code Exchange (PKCE) by OAuth Public Clients specification.
	When enabled, this client must include a one-time string value through the use of the code_challenge parameter in its authorization request. For more information, see . It must also submit the corresponding code verifier through the code_verifier parameter in its token request when exchanging an authorization code for an access token. For more information, see .
	This check box is not selected by default.
Polling Interval (seconds)	Specifies the number of seconds that the client must wait between its attempts to check for the authorization results at the token endpoint. When PingFederate receives a token request within this time interval, it returns a slow_down error message to the client.
	A valid value ranges from 1 to 3600.
	The default value is 3.
Policy	Specifies the CIBA request policy associated with the client.
	PingFederate uses CIBA request policies to determine various aspects of CIBA authentication request, such as the maximum lifetime of authentication requests, the validity of unsigned login hint tokens, and the mapping configuration of identity hints.
	Select an existing CIBA policy. You may also leave the selection of Default to indicate that PingFederate should use the CIBA request policy that has been designated as the default CIBA request policy in Applications # OAuth # CIBA Request Policies .
Require CIBA Signed Requests	Determines whether the client must transmit request parameters in a single, self-contained parameter. The parameter name is request. The value of the request parameter is a signed JWT whose claims represent the request parameters of the authorization request. The OpenID Connect specification calls this JWT a request object.
	This check box is not selected by default.
	Note that if client-initiated backhannel authentication (CIBA) signed requests are required, the dynamic client registration must include either the JWKS URL or the actual JWKS.
token exchange	This field determines which token exchange processor policy PingFederate uses when the OAuth server receives an OAuth token exchange request from the client. If you select Default , PingFederate uses the token exchange processor policy that was set as the default on the Token Exchange Processor Policy Management tab, under Applications # Token Exchange # Processor Policies .
	For more information, see <i>OAuth token exchange</i> .

Selecting client registration policies

Client registration policies can provide additional control over which registrations and configurations are accepted and stored for each client created with the OAuth 2.0 Dynamic Client Registration protocol.

About this task

If multiple policies are configured, PingFederate executes all of them based on the display order. If PingFederate completes the current policy, it moves on to the next policy. Otherwise, PingFederate returns an error message to the developers.



Note:

PingFederate must complete all policies successfully before a client can be created with the OAuth 2.0 Dynamic Client Registration protocol.

Steps

- 1. Go to System # OAuth Settings # Client Registration Policies.
- 2. Optional: Select a Client Registration Policy instance from the **Available Policies** list and click **Add**.



Important:

Select this option to add a layer of protection against unwanted client registrations.

If you have not yet defined the desired Client Registration Policy instance, click Manage Client Registration Policies.

3. Optional: Repeat the previous step to add other Client Registration Policy instances.

Add as many Client Registration Policy instances as necessary. Click the up and down arrows to adjust the execution order. Use the **Delete** and **Undelete** buttons to remove an existing instance or cancel the removal request.

Reviewing client settings

On the **Summary** tab, review your client settings.

Steps

- To amend your configuration, click the corresponding tab title and then follow the steps to complete the task.
- To keep your changes, click Save.
- To discard your changes, click Cancel.

Managing Client Registration Policy instances

The Client Registration Policy plugin allows you to write custom processing rules to provide additional control over which registrations and configurations are accepted and stored for each client created with the OAuth 2.0 Dynamic Client Registration protocol.

About this task

Depending on the technical requirements of your use cases, you can create Client Registration Policy plugins using the PingFederate SDK. After deploying your plugins, you can create and configure instances of them. Configuration requirements vary based on your custom solutions. When you are ready to configure dynamic client registration, add your policies to its configuration.

Steps

1. Implement the DynamicClientRegistrationPlugin interface.

For more information, refer to the Javadoc for the DynamicClientRegistrationPlugin interface, the SoftwareStatementValidatorPlugin.java file for a sample implementation, and the SDK developer's guide for build and deployment information.



Tip:

The Javadoc for PingFederate and the sample implementation are in the <pf install>/ pingfederate/sdk directory.

2. Create, modify, or remove one or more instances.

Choose from:

- To configure a new instance, click Create New Instance.
- To modify an existing instance, select it under Instance Name.
- To remove an existing instance or to cancel the removal request, click Delete or Undelete under Action.



Note:

You can remove a Client Registration Policy instance only if it is not currently in-use by dynamic client registration.

To save the plugin configuration, click Save.

Result



Important:

A Client Registration Policy instance is not enforced, or executed as part of the dynamic client registration process, until it is selected on the Client Registration Policies window.

Configuring a Client Registration Policy instance

For additional control over clients created with the OAuth 2.0 Dynamic Client Registration protocol, use the Client Registration Policies window to create or modify a Client Registration Policy instance.

Steps

- 1. Go to System # OAuth Settings # Client Registration Policies.
 - Choose from:
 - To configure a new instance, click Create New Instance.
 - To modify an existing instance, select it under Instance Name.
- 2. On the **Type** tab, enter a name and an ID for a new instance, and then select a plugin from the list.

When modifying an existing policy plugin instance, you can only change the Instance Name field.

If the **Type** list does not contain the desired Client Registration Policy plugin, create one using the PingFederate SDK. For more information, refer to the Javadoc for the DynamicClientRegistrationPlugin interface, the

SoftwareStatementValidatorPlugin.java file for a sample implementation, and the SDK developer's guide for build and deployment information.



The Javadoc for PingFederate and the sample implementation are in the <pf install>/ pingfederate/sdk directory.

3. In the Instance Configuration tab, follow the on-screen instructions to configure the Client Registration Policy instance.



Note:

This window varies depending on the selected Client Registration Policy plugin.

- 4. On the **Summary** tab, review the plugin configuration. Click **Done**.
- 5. In the Client Registration Policies window, click Save.

Result



Important:

A Client Registration Policy instance is not enforced, or executed as part of the dynamic client registration process, until it is selected on the Client Registration Policies window.

Configuring a Response Type Constraints instance

The Response Type Constraints policy plugin allows administrators to control which flows are allowed for clients created through the OAuth 2.0 Dynamic Client Registration protocol.

About this task

Configure an instance of the Response Type Constraints policy to limit which of the following response types parameter values are allowed:

- code
- code id_token
- code id_token token
- code token
- id token
- id token token
- token

For more information about flows and response types, see the OpenID Connect specification.

Steps

- 1. Go to System # OAuth Settings # Client Registration Policies. Choose from:
 - To configure a new instance, click Create New Instance.
 - To modify an existing instance, select it under Instance Name.
- 2. On the **Type** tab, enter a name and an ID for a new instance, and then select **Response Type** Constraints from the Type list.

When modifying an existing policy plugin instance, you can only change the Instance Name field.

3. On the Instance Configuration tab, clear the applicable check boxes to remove the unwanted response types.



Note:

All response types are allowed by default.

- 4. On the **Summary** tab, review the plugin configuration. Click **Done**.
- 5. In the Client Registration Policy Instances window, click Save.

Result



Important:

Like other Client Registration Policy plugins, an instance of the Response Type Constraints policy plugin is not enforced, or executed as part of the dynamic client registration process, until it is selected in **System** # OAuth Settings # Client Registration Policies. If it is selected in the Client Registration Policies window, PingFederate discards all restricted response types when processing client registrations. If no response type is allowed, PingFederate rejects the registration and returns an error message to the originator.

Managing OAuth clients

An OAuth client application interacts with an OAuth authorization server to obtain the required access tokens to call OAuth-protected services at the resource server.

About this task

The Clients window displays 20 clients at a time. You can sort the display order by name or ID. You can use the pagination controls to navigate through the rest of the clients or search clients by name or ID. A client is included in the search results if its name or ID is a partial, case-insensitive match to the search term.

Steps

To manage OAuth clients, go to Applications # OAuth # Clients.

Action	Steps
To add a client	Click Add Client and complete the configuration in the Client window.
To edit a recently modified client	Select the client and update the configuration in the Client window.
To enable or disable one or more clients	Click their toggle switches and then click Save .
To remove a client or cancel the removal request	Use the Delete and Undelete buttons for the applicable client and then click Save .

Result

PingFederate stores client records in XML files by default. On-disk storage allows you to manage clients using the administrative console and the administrative API. Client records are part of the configuration archive.

Alternatively, you can configure PingFederate to store client records externally, which allows you to manage client records through the OAuth Client Management Service or enable dynamic client registration for your partner-developers. In this case, client records are not part of the configuration archive. Instead, PingFederate stores them on a database server, a directory server, or another storage medium through the use of the PingFederate SDK. For more information, see *OAuth client datastores* on page 895.

Configuring an OAuth client

The **Clients** window provides controls over the usage and behavior of the applications requesting access to protected resources through the PingFederate OAuth authorization server.

Steps

1. Go to **Applications** # **OAuth** # **Clients** and configure the OAuth client to suit your use cases. The following table describes each field.

Field	Description
Client ID	A unique identifier the client provides to the resource server (RS) to identify
(Required)	itself. This identifier is included with every request the client makes.
Name	A descriptive name for the client instance. This name appears when the
(Required)	user is prompted for authorization.
	Tip:
	If you want to localize the displayed name, you can enter a unique alias here, then use the same alias in language resource files.
Description	A description of what the client application does. This description appears when the user is prompted for authorization.
	Tip:
	If you want to localize the displayed description, you can enter a unique alias here, then use the same alias in language resource files.

Client Authentication

The authentication method that the client uses.

None

Select this option if your use case does not require client authentication. This is the default selection.



Note:

A selection other than **None** is required for any of the following use cases:

- This client uses the Client Credentials grant type. See the Allowed Grant Types check boxes.
- This client signs its ID tokens using an HMAC signing algorithm. See the ID Token Signing Algorithm field.
- This client is allowed to access the Session Revocation API. See the Allow Access to Session Revocation API check box.
- This client is allowed to access the Session Management API. See the Allow Access to Session Management API check box.

Client Secret

Select this option for HTTP Basic authentication.

- To create a strong, random alphanumeric string or to manually enter a secret, click Generate Secret.
- To modify an existing secret, select the **Change Secret** check box. Then, click **Generate Secret** to create a strong random alphanumeric string or manually enter a secret.

Client TLS Certificate

Select this option for mutual TLS certificate-based authentication; recommended for client applications where security policies prohibit storing passwords.

- Select a trusted CA from the Issuer list.
 - These are CA certificates imported into PingFederate. You can review them on the Security # Certificate & Key Management # Trusted CAs window. Alternatively, you can select Trust Any to trust all the issuers found in the list.
- Enter the client-certificate subject DN in the **Subject DN** or extract the subject DN from the certificate if the certificate is stored on an accessible file system.



Important:

pf.secondary.https.port in the table under If choosing this option, you must configure a secondary PingFederate HTTPS port. See the property Configuring PingFederate properties on page 752.

Private Key JWT

If choosing this option, you must configure a secondarySelect this option for the private_key_jwt client authentication method, as defined in Client Authentication in the OpenID Connect specification (openid.net/specs/openid-connectcore-1_0.html#ClientAuthentication).

Copyright ©2024

Field

Description

Require Signed Requests

Determines whether the client must transmit request parameters in a single, self-contained parameter. The parameter name is request. The value of the request parameter is a signed JWT whose claims represent the request parameters of the authorization request. The OpenID Connect specification calls this JWT a request object.



Note:

If a client includes in an authorization request a request parameter, other than client id and response type, as a parameter outside of the signed request object and a claim inside of the signed request object, PingFederate always uses the claim value found inside the signed request object to process the request further.

For the client id and response type request parameters, the values outside of the signed request object must match the claim values inside of the signed request object. If the values do not match, PingFederate returns an error message to the client.

If a request parameter is found only outside of the signed request object, PingFederate ignores the request parameter and returns no error message.



Per OAuth and OpenID Connect specifications, a client must always include in an authorization request the client id, response type, and scope request parameters outside of the signed request object.

For more information about request objects, see OpenID Connect specifications.

Request Object Signing Algorithm

The signing algorithm that the client must use to sign its request objects for transmission of request parameters.

Applicable only when the client may send its authorization requests using request objects.

If PingFederate is either deployed to run in a Java 11 runtime environment or integrated with a hardware security module (HSM) and configured to use static keys for OAuth and OpenID Connect, additional RSASSA-PSS signing algorithms become available for selection. For more information on HSM integration and static keys, see and .

The default selection is Allow Any, which allows the client to use any of the signing algorithms from the list.

Field	Description
JWKS URL and JWKS	The URL of the JSON Web Key Set (JWKS) or the actual JWKS from the client.
	If the client is configured to use the private_key_jwt client authentication method, to transmit request parameters in signed request objects, or to transmit CIBA request parameters in signed request objects, only one of the previous values is required for PingFederate to verify the authenticity of the JWTs.
	Either value can be defined even if the client is not configured to use JWTs for authentication or transmission of request parameters. This flexibility allows the client to transmit request parameters in signed request objects for some requests and without the use of signed request objects for some other transactions. For information on runtime processing, see .
	If the client signs its JWTs using an RSASSA-PSS signing algorithm, PingFederate must be deployed to run in a Java 11 runtime environment, or be integrated with a hardware security module (HSM) to process the digital signatures.For more information on HSM integration, see .
	If the client is configured to encrypt ID tokens using an asymmetric encryption algorithm, either the JWKS URL or the actual JWKS must be provided. See the ID Token Key Management Encryption Algorithm setting.
Redirection URIs	URIs where the OAuth AS may redirect the resource owner's user agent after authorization is obtained. The authorization code and implicit grant types require at least one redirection URI.
	Enter a fully qualified URL and click Add for each entry required. Wildcards are allowed. However, for security reasons, make the URL as restrictive as possible, for example, https://www.example.com/OAuthClientApp/callback.jsp.
	Important:
	If more than one URI is added or if a single URI uses wildcards, then the authorization code grant and the token requests must contain a specific matching redirect_uriparameter when contacting the authorization endpoint (/as/authorization.oauth2) and token endpoint (/as/token.oauth2).
Logo URL	The location of the logo used on user-facing OAuth grant authorization and revocation pages. For best results with the installed HTML templates, the recommended size is 72 x 72 pixels.

Field Description **Allow Authentication** When selected, the client can initiate an authentication API OAuth flow **API OAuth Initiation** through the authorization endpoint without needing to handle HTTP redirections. When enabling this feature, consider the following: Redirection URLs are optional, but without a redirection URL, browserbased OAuth flows will not work. This flow does not support the user-facing scope consent page, Request for Approval. So, enabling this feature automatically enables the Bypass Authorization Approval feature and Restrict Common Scopes feature. The client must manage the PF cookie and, if persistent authentication sessions are configured, the PF.PERSISTENT cookie. For more information, see Mobile application authentication through REST APIs on page 1159. **Bypass Authorization** When selected, resource-owner approval for client access is assumed, and **Approval** PingFederate no longer presents to the user an authorization consent page or redirects to a trusted web application that is responsible to prompt the user for authorization for this client. For instance, use this setting when you want to deploy a trusted application and authenticate end users with an identity provider (IdP) adapter or IdP connection. **Restrict Common** Controls whether all existing and future common scopes and scope groups **Scopes** should be made available to the client, or only the select few. When selected, the administrative console displays a list of existing common scopes and scope groups. Choose the common scope and scope groups that are intended for the client. The rest and any common scopes and scope groups created in the future become invalid for the client. In other words, if the client tries to use such scope or scope group, it will receive an invalid scope error message from PingFederate. When cleared, all existing common scopes and scope groups and those created in the future are available to the client. This is the default behavior. 1 Note: Depending on the configured dynamic scope patterns and whether they are defined as common or exclusive dynamic scopes, this setting can impact the results of scope evaluation. The default scope, however, is always allowed for and available to all clients. For detailed information, see the Dynamic scope evaluation and per-client scope management section in Scopes and scope management on page 486.

Field Description **Exclusive Scopes** Controls whether any exclusive scopes and scope groups should be made available to the client. When selected, the administrative console displays a list of existing exclusive scopes and scope groups. Choose the exclusive scopes and scope groups that are intended for the client. The rest and any exclusive scopes and scope groups created in the future become invalid for the client. In other words, if the client tries to use such scope or scope group, it will receive an invalid scope error message from PingFederate. When cleared, no exclusive scopes and scope groups are available to the client. This is the default behavior. Note: Depending on the configured dynamic scope patterns and whether they are defined as common or exclusive dynamic scopes, this setting can impact the results of scope evaluation. The default scope, however, is always allowed for and available to all clients. For detailed information, see the Dynamic scope evaluation and per-client scope management section in Scopes and scope management on page 486. Allowed Grant Types The grant types that this client can use. Select at least one of the following: Authorization Code Implicit Refresh Token **Client Credentials Device Authorization Grant** CIBA **Token Exchange Resource Owner Password Credentials Assertion Grants**

Access Token Validation (Client is a Resource Server)

For more information about each grant type, see *Grant types*.

There is no default selection.

Field	Description		
Restrict Response Types	Select this check box to limit the response_type parameter values that this client can use.		
	Available response types are:		
	 code code id_token code id_token token code token id_token id_token token token 		
	For more information about these response types, see Definitions of Multiple-Valued Response Type Combinations.		
	The Restrict Response Type check box is not selected by default. If selected, you must select at least one allowable response_type parameter value.		
	Additionally, the Restricted Response Types and Allowed Grant Types settings must be configured in tandem because certain response types require one or more grant types, and vice versa. The following table describes their relationship.		
	Response type	Grant types	
	code	Authorization Code	
	code id_token	Authorization Code and Implicit	
	code id_token token	Authorization Code and Implicit	
	code token	Authorization Code and Implicit	
	id_token	Implicit	
	id_token token	Implicit	
	token	Implicit	
Default Access Token Manager	Determines the default Access Toker client.	n Management (ATM) instance for this	
Validate Against All	Applicable only to resource server clients.		
Eligible Access Token Managers	If selected, this resource server client is not required to specify the additional access_token_manager_id or aud parameters to disambiguate the ATM instance in its token validation requests. When the resource server client does not specify the desired ATM instance, PingFederate validates the access tokens against all eligible ATM instances. This simplifies interactions with PingAccess by avoiding the need to align resource URIs between PingAccess and PingFederate.		
	This check box is not selected by default.		

Field Require Proof Key for Code Exchange (PKCE)

Description

Displayed only when the client is configured to support the authorization code grant type.

Determines whether the client must provide certain parameters to reduce the risk of authorization code interception attack. For more information, see the Proof Key for Code Exchange (PKCE) by OAuth Public Clients specification.

When enabled, this client must include a one-time string value through the use of the code challenge parameter in its authorization request. For more information, see . It must also submit the corresponding code verifier through the code verifier parameter in its token request when exchanging an authorization code for an access token. For more information, see .

This check box is not selected by default.

Lifetime

Persistent Grants Max Overrides the Persistent Grant Max Lifetime value set globally in System # OAuth Settings # Authorization Server Settings.

Select one of the following options:

- Use Global Setting (default)
- **Grants Do Not Expire**
- A custom value in days, hours, or minutes.



This setting can be overridden per grant-mapping configuration through the use of an extended persistent grant attribute PERSISTENT GRANT LIFETIME. The PERSISTENT GRANT LIFETIME attribute is defined in System # OAuth Settings # Authorization Server **Settings**. When this attribute is active, you can set the lifetime of persistent grants based on the outcome of attribute mapping expressions in individual grant-mapping configurations. For grant-mapping configurations that do not require this fine-grain control, you can configure them to use the default value.

Persistent Grants Idle Timeout

Overrides the Persistent Grant Idle Timeout field value set globally in System # OAuth Settings # Authorization Server Settings.

Select one of the following options:

- Use Global Setting (default)
- **Grants Do Not Timeout Due To Inactivity**
- A custom value in days, hours, or minutes.

If you configure an idle timeout value, the idle timeout window slides when a persistent grant updates. When you have an idle timeout value configured without a maximum lifetime, persistent grants remain valid until they expire due to inactivity or until the grant storage revokes or removes them. When you have an idle timeout value configured with a maximum lifetime, persistent grants remain valid until they expire due to inactivity or lifetime expiration or until the grant storage removes them. For more information, see *Transient grants and persistent grants* on page 79.

Field	Description
Refresh Token Rolling Policy	Overrides the Roll Refresh Token Values setting configured globally in System # OAuth Settings # Authorization Server Settings.
	Select one of the following options:
	Use Global Setting (default)Roll
	This value does not override the Minimum Interval to Roll Refresh Tokens (Hours) value set in the Authorization Server Settings window. - Don't Roll

Field Description

OpenID Connect

ID Token Signing Algorithm

Select the signing algorithm for the ID tokens from the list. The default algorithm is RSA using SHA-256.

If PingFederate is either deployed to run in a Java 11 runtime environment or integrated with a hardware security module (HSM) and configured to use static keys for OAuth and OpenID Connect, additional RSASSA-PSS signing algorithms become available for selection. For more information on HSM integration and static keys, see and.



Note:

If static keys for OAuth and OpenID Connect are enabled, Ellipticcurve cryptography (EC) algorithms that have not been configured with an active static keys are hidden.

Changes made in the static-key configuration might affect runtime transactions and require additional changes here. For more information, see .

ID Token Key Management Encryption Algorithm

The algorithm used to encrypt or otherwise determine the value of the content encryption key.

PingFederate supports symmetric algorithms, such as Direct Encryption with symmetric key, AES ... Key Wrap. and AES-GCM ... key encryption, and asymmetric algorithms, such as ECDH-ES, ECDH-ES ... Key Wrap, and RSAES OAEP.

ID Token Content Encryption Algorithm

The content encryption algorithm used to perform authenticated encryption on the plain text payload of the token.

Required if an algorithm is selected from the ID Token Key Management Encryption Algorithm list.

Policy

Select a specific OpenID Connect policy from the list.



Note:

If the Track User Sessions for Logout check box is selected in the System # OAuth Settings # Authorization Server Settings window, PingFederate also displays the PingAccess Logout Capable and Logout URIs fields.

PingAccess Logout Capable

When selected, PingFederate sends logout requests through the browser to an OpenID Connect endpoint in PingAccess as part of the logout process. For more information, see *OpenID Connect endpoints*. This check box is not selected by default.

Logout URIs

Enter additional endpoints at the relying parties as needed. PingFederate sends requests to these URIs through the browser as part of the logout process. The relying parties must return an image in their logout responses, otherwise PingFederate returns an error

For PingFederate to skip this confirmation step, select this check box.

This check box is not selected by default.

Field	Description
CIBA	Displayed only if the CIBA grant type is enabled for the client.
	Token Delivery Method
	The token delivery method that the client supports. PingFederate supports poll and ping.
	Select Poll if the client can check for the authorization results at the token endpoint periodically. Select Ping if the client prefers to wait for a ping callback message from PingFederate as a signal that the authorization result is ready for pickup.
	The default selection is Poll .
	Notification Endpoint
	The client's notification endpoint, to which PingFederate sends its ping call back messages.
	Required and displayed only if ping is the configured token delivery method.
	Polling Interval (seconds)
	Specifies the number of seconds that the client must wait between its attempts to check for the authorization results at the token endpoint. When PingFederate receives a token request within this time interval, it returns a slow_down error message to the client.
	A valid value ranges from 1 to 3600.
	The default value is 3.
	Policy
	Specifies the CIBA request policy associated with the client. PingFederate uses CIBA request policies to determine various aspects of CIBA authentication request, such as the maximum lifetime of authentication requests, the validity of unsigned login hint tokens, and the mapping configuration of identity hints.
	Select an existing CIBA policy. You may also leave the selection of Default to indicate that PingFederate should use the CIBA request policy that has been designated as the default CIBA request policy in Applications # OAuth # CIBA Request Policies .
	User Code Support
	Indicates whether the client supports user code.
	The purpose of this code is to authorize the transmission of an authentication request to the user's authentication device.
	This check box is not selected by default.
	When user code support is enabled, the associated CIBA request policy must also be user code enabled.
	Require CIBA Signed Requests
	Determines whether the client must transmit request parameters in a single, self-contained parameter. The parameter name is request. The value of the request parameter is a signed JWT whose claims represent the request parameters of the authorization request. The OpenID Connect specification calls this JWT a request object.

This check box is not selected by default.

If CIBA signed requests are required, the client must also be configured with either the JWKS URL or the actual JWKS from the $\,$

- 2. To enable or disable the client, click the toggle switch.
- 3. Optional: On the **System** tab, add, remove, or update one or more values for any extended properties defined in **System** # **Server** # **Extended Properties**.

Extended property values can serve as metadata. They can also help drive authentication requirements. For more information, see *Extended properties* on page 856.

4. Click Save.

Grant contract mapping

In the first stage of the OAuth attribute mapping process, PingFederate maps attributes to persistent grant contracts.

You configure PingFederate to use attributes from the following sources:

- authentication policy contracts
- authentication sources, such as IdP adapter instances and IdP connections
- password credential validator instances for resource owner credentials

Depending on the attribute source, use one of the following PingFederate windows to configure the grant contract mapping:

- Authentication # OAuth # Policy Contract Grant Mapping
- Authentication # OAuth # IdP Adapter Grant Mapping
- Authentication # OAuth # Resource Owner Credentials Grant Mapping

These windows also let you configure issuance criteria to control whether PingFederate fulfills the contract.

Persistent grants, and any associated attributes and their values, remain valid until the grants expire or until PingFederate explicitly revokes them or cleans them up.

For more information about the OAuth attribute mapping process, see *Mapping OAuth attributes* on page 81

Managing IdP adapter grant mapping

Use the **IdP Adapter Grant Mapping** to map authentication source values into persistent grants. Persistent grants and any associated attributes and their values remain valid until the grants expire or until PingFederate explicitly revokes or cleans them up.

About this task

The USER_KEY attribute is the identifier of the persistent grants. The USER_NAME attribute presents the name shown to the resource owner on OAuth user-facing pages. If extended attributes are defined in System # OAuth Settings # Authorization Server Settings, configure a mapping for each attribute. You can optionally set up datastore queries to supplement values returned from the source. This mapping configuration is suitable for the Authorization Code and Implicit grant types.

Steps

Go to Authentication # OAuth # IdP Adapter Grant Mapping and perform one of the following actions.

Configuring IdP adapter attribute sources and user lookup

You can optionally set up datastore gueries to supplement values returned from the identity provider (IdP) adapter attribute source. Once the configuration is complete, you can fulfill a contract or verify a condition in the Token Authorization framework using the results from the queries.

Steps

- To set up datastore queries, click Add Attribute Source.
 - Follow the Attribute Sources & User Lookup window to complete the setup. For configuration steps, see Datastore query configuration on page 362.
- To skip this option, click Next.

Fulfilling IdP adapter grant mapping

On the Contract Fulfillment tab, map authentication source values into persistent grants. Persistent grants and any associated attributes and their values remain valid until the grants expire or until PingFederate explicitly revokes or cleans them up.

About this task

The USER KEY attribute is the identifier of the persistent grants. The USER NAME attribute presents the name shown to the resource owner on OAuth user-facing pages. If extended attributes are defined in System # OAuth Settings # Authorization Server Settings, configure a mapping for each attribute.



Important:

The USER KEY attribute values must be unique across all end users, because the USER KEY attribute is the user identifier to store and to retrieve persistent grants. For example, the samaccountName attribute value of an end user in one domain might match that of another end user in another domain. In this case, you can map the Subject DN attribute to the USER KEY attribute.

Steps

1. Go to Authentication # OAuth # IdP Adapter Grant Mapping and select your mapping, or click Add Mapping.

2. On the Contract Fulfillment tab, select a source from the Source list and then select or enter a value for each attribute in the contract.

You can map each attribute from one of the following sources:

• Adapter

When selected, the associated Value drop-down list contains attributes configured in the IdP adapter instance.

Context

Values are returned from the context of the transaction at runtime.



Note:

If PERSISTENT GRANT LIFETIME is an extended attribute in System # OAuth Settings # Authorization Server Settings, you can set the lifetime of persistent grants based on the outcome of attribute mapping expressions, or the per-client Persistent Grants Max Lifetime setting.

- To set lifetime based on the per-client Persistent Grants Max Lifetime setting, select Context from the Source list and Default Persistent Grant Lifetime from the Value list.
- To set lifetime based on the outcome of attribute mapping expressions, select **Expression** as the source and enter an OGNL expression in the Value field.

If the expression returns a positive integer, the value represents the lifetime of the persistent grant in minutes.

If the expression returns the integer 0, PingFederate does not store the grant and does not issue a refresh token.

If the expression returns any other value, PingFederate sets the lifetime of the persistent grant based on the per-client Persistent Grants Max Lifetime setting.

To set a static lifetime, select **Text** from the **Source** list and enter a static value in the **Value** field.

This is suitable for testing purposes, or cases where the persistent grant lifetime must always be set to a specific value.

As the HTTP Request context value is retrieved as a Java object rather than text, OGNL expressions are ideal to evaluate and return values.

• Extended Client Metadata

Values are returned from the client record.

• LDAP/JDBC/Other

Values are returned from your datastore, if used.

• Expression

If enabled, this option provides more complex mapping capabilities, such as transforming incoming values into different formats. All of the variables available for text entries are available for expressions.

• No Mapping

This option ignores the Value field.

Text

You can enter text only, or mix text with references to the attributes returned from the adapter instance, using the \${attribute} syntax.

You can also enter values from your datastore using the \${ds.attribute} syntax, where attribute is any of the datastore attributes you have selected.

Defining issuance criteria for OAuth IdP adapter mapping

Individual attributes within policy contracts can further determine whether PingFederate approves or rejects requests. You can define those criteria to satisfy or you can choose to skip this configuration.

About this task

On the Issuance Criteria tab, define the criteria to satisfy in order for PingFederate to further process a request. Use this token authorization feature to conditionally approve or reject requests based on individual attributes.

Begin this optional configuration by choosing the source that contains the attribute to verify. Some sources are common to almost all use cases, such as Mapped Attributes. Other sources depend on the type of configuration, such as JDBC. Irrelevant sources are automatically hidden. Once you select a source, choose the attribute to verify. Depending on the selected source, the available attributes or properties vary. Specify the comparison condition and the desired value to compare to.

You can define multiple criteria, which must all be satisfied in order for PingFederate to move a request to the next phase. A criterion is satisfied when the runtime value of the selected attribute matches or does not match the specified value, depending on the chosen comparison method. The multi-value contains ... or multi-value does not contain ... comparison methods are intended for attributes that can contain multiple values. Such a criterion is considered satisfied if one of the multiple values match or does not match the specified value. Values are compared verbatim. If you require complex evaluations, including conditional criteria or partial matching, define them using attribute mapping expressions.



Note:

All criteria defined must be satisfied, or evaluated as true, for a request to move forward, regardless of how the criteria were defined. As soon as one criterion fails, PingFederate rejects the request and returns an error message.

Steps

- 1. Go to Authentication # OAuth # IdP Adapter Grant Mapping. Select an adapter instance, and then click Issuance Criteria.
- 2. From the **Source** list, select the attribute's source.

Depending on the selection, the **Attribute Name** list populates with associated attributes. See the following table for more information.

Source	Description
Adapter	Select to evaluate attributes from the IdP adapter instance.
Context	Select to evaluate properties returned from the context of the transaction at runtime.
	Note: As the HTTP Request context value is retrieved as a Java object rather than text, attribute mapping expressions are more appropriate to evaluate and return values.
Extended Client Metadata	Select to evaluate OAuth client metadata.

Source	Description
JDBC, LDAP, or other types of datastore (if configured)	Select to evaluate attributes returned from a data source.
Mapped Attributes	Select to evaluate the mapped attributes.

- 3. From the **Attribute Name** list, select the attribute to be evaluated.
- 4. From the **Condition** list, select the comparison method.

Available methods:

- equal to
- equal to (case insensitive)
- equal to DN
- not equal to
- not equal to (case insensitive)
- not equal to DN
- multi-value contains
- multi-value contains (case insensitive)
- multi-value contains DN
- multi-value does not contain
- multi-value does not contain (case insensitive)
- multi-value does not contain DN



The first six conditions are intended for single-value attributes. Use one of the multi-value ... conditions for PingFederate to validate whether one of the attribute values matches the specified value. When an attribute has multiple values, using a single-value condition causes the criteria to fail.

5. In the **Value** field, enter the comparison value.



Note:

Values are compared verbatim. If you require complex evaluations, including conditional criteria or partial matching, define them using attribute mapping expressions. For more information, see Attribute mapping expressions on page 214.

6. In the **Error Result** field, enter a custom error message.

This value is used by the error description protocol field. Using an error code in the Error Result field allows an application to process the code in several ways, such as displaying an error message or e-mailing an administrator.

To use localized descriptions, enter a unique alias in the Error Result field, such as someIssuanceCriterionFailed. Insert the same alias with the desired localized text in the applicable language resource files, located in the <pf install>/pingfederate/server/ default/conf/language-packs directory.

If not defined, PingFederate returns ACCESS DENIED when the criterion fails at runtime.

- 7. Click Add.
- 8. Optional: Repeat to add more criteria.

- 9. If you require complex evaluations, including conditional criteria or partial matching, define them using attribute mapping expressions. For more information, see Attribute mapping expressions on page 214.
 - a. Click Show Advanced Criteria.
 - b. In the **Expression** field, enter the required expressions.
 - c. Optional: In the Error Result field, enter an error code or message.



Note:

If the expressions resolve to a string value instead of true or false, the returned value overrides the Error Result field value.

- d. Click Add.
- e. Optional: Click **Test**, enter values in the applicable fields, and verify the results.
- f. Optional: Repeat to add multiple criteria using attribute mapping expressions.

Reviewing the IdP adapter mapping

On the **Summary** tab, review your identity provider (IdP) adapter mapping.

Steps

- To amend your configuration, click the corresponding tab title and then follow the steps to complete the task.
- To keep your changes, click Save.
- To discard your changes, click Cancel.

Configuring IdP connection grant mapping

Use this configuration to map values obtained from the single sign-on (SSO) tokens into the persistent grants. Persistent grants remain valid until the grant expires or is explicitly revoked.

About this task

The USER KEY attribute is the identifier of the persistent grants. The USER NAME attribute presents the name shown to the resource owner on OAuth user-facing pages. If extended attributes are defined in System # OAuth Settings # Authorization Server Settings, configure a mapping for each attribute. You can optionally set up datastore gueries to supplement values returned from the source. This mapping configuration is suitable for the Authorization Code and Implicit grant types.

Steps

- 1. Go to Authentication # Integration # IdP Connections and select an existing identity provider (IdP) connection or click Create Connection.
- 2. On the Connection Type tab, select the Browser SSO Profiles check box and the applicable protocol.
- 3. On the Connection Options window, select the Browser SSO check box and then select the OAuth Attribute Mapping check box.



You can also select other options on the **Connection Type** and **Connection Options** tabs. If you do, you will be prompted to complete the required configuration. For simplicity, this topic only focuses on the **OAuth Attribute Mapping** configuration.

4. On the **General Info** tab, enter the required information.

- 5. On the Browser SSO tab, click Configure Browser SSO and follow the steps to complete the User-Session Creation tab.
- 6. On the OAuth Attribute Mapping tab, select the Map directly into Persistent Grant option, and then click Configure OAuth Attribute Mapping to continue.

Alternatively, if you have mapped an authentication policy contract (APC) in **User-Session Creation** # Target Session Mapping, you can select the Map to OAuth via Authentication Policy Contract option, and then select the applicable APC from the list.

Choosing an OAuth datastore

You can optionally set up OAuth datastore queries to supplement values returned from the source.

Steps

On the **Data Store** tab, perform one of the following actions.

Choose from:

- To set up datastore queries, select a datastore from the Active Data Store list and then click Next. For configuration steps, see *Datastore guery configuration* on page 362.
- To skip this optional configuration, select No Data Store, and then click Next.

Fulfilling OAuth attribute mapping

On the **Contract Fulfillment** tab, map authentication source values into persistent grants.

About this task

The USER KEY attribute is the identifier of the persistent grants. The USER NAME attribute presents the name shown to the resource owner on OAuth user-facing pages. If extended attributes are defined in System # OAuth Settings # Authorization Server Settings, configure a mapping for each attribute.



Important:

The USER KEY attribute values must be unique across all end users because the USER KEY attribute is the user identifier to store and to retrieve persistent grants. For example, if you are configuring an OAuth Attribute Mapping configuration on a SAML 2.0 IdP connection and the SAML SUBJECT attribute uniquely identifies all end users, you can map the SAML SUBJECT attribute to the USER KEY attribute.

Steps

- 1. For each attribute, select a source from the list and then choose or enter a value.
 - AccountLink

When selected, the Value list is populated with Local User ID. You can map Local User ID to an attribute that represents the user identifier, such as the USER KEY attribute. This source appears

only if you have elected to use account linking for a target session on the Identity Mapping window.

Assertion or Provider Claims

When selected, the Value list is populated with attributes from the SSO token. Select the desired attribute from the list.

For example, to map the value of **SAML SUBJECT** from a SAML assertion as the value of the USER KEY user identifier on the contract, select Assertion from the Source list and SAML SUBJECT from the Value list.

Context

When selected, the Value list populates with the available context of the transaction. Select the desired context from the list.



Note:

As the HTTP Request context value is retrieved as a Java object rather than text, use OGNL expressions to evaluate and return values.



Note:

If you are configuring an OAuth Attribute Mapping configuration and have added PERSISTENT GRANT LIFETIME as an extended attribute in the Authorization Server Settings window, you can set the lifetime of persistent grants based on the outcome of attribute mapping expressions or the per-client Persistent Grants Max Lifetime setting.

- To set lifetime based on the per-client Persistent Grants Max Lifetime setting, select Context from the Source list and Default Persistent Grant Lifetime from the Value list.
- To set lifetime based on the outcome of attribute mapping expressions, select Expression as the source and enter an OGNL expression in the Value field.

If the expression returns a positive integer, the value represents the lifetime of the persistent grant in minutes.

If the expression returns the integer 0, PingFederate does not store the grant and does not issue a refresh token.

If the expression returns any other value, PingFederate sets the lifetime of the persistent grant based on the per-client Persistent Grants Max Lifetime setting.

To set a static lifetime, select **Text** from the **Source** list and enter a static value in the **Value** field.

This is suitable for testing purposes, or cases where the persistent grant lifetime must always be set to a specific value.

Extended Client Metadata

Values are returned from the client record.

• LDAP, JDBC, or Other

When selected, the Value list is populated with attributes selected from the datastore. Select the desired attribute from the list.

Expression

When enabled, this option provides more complex mapping capabilities, such as transforming incoming values into different formats. Select Expression from the Source list, click Edit under Actions, and compose your OGNL expressions. All variables available for text entries are also available for expressions. For more information, see **Text**.

Expressions are not enabled by default. For more information about enabling and editing OGNL expressions, see .

No Mapping

When selected, no value selection is necessary.

Text

When selected, the text you enter is used at runtime. You can mix text with references to any of the values from the SSO token, using the \${attribute} syntax.

When applicable, you can also enter values from your datastore using the \${ds.attribute} syntax, where attribute is any attribute that you have selected from the datastore.

2. Click Next.

Defining issuance criteria for OAuth attribute mapping

Individual attributes within policy contracts can further determine whether PingFederate approves or rejects requests. You can define those criteria to satisfy or you can choose to skip this configuration.

About this task

On the **Issuance Criteria** tab, define the criteria to satisfy in order for PingFederate to further process a request. Use this token authorization feature to conditionally approve or reject requests based on individual attributes.

Begin this optional configuration by choosing the source that contains the attribute to verify. Some sources are common to almost all use cases, such as Mapped Attributes. Other sources depend on the type of configuration, such as JDBC. Irrelevant sources are automatically hidden. Once you select a source, choose the attribute to verify. Depending on the selected source, the available attributes or properties vary. Specify the comparison condition and the desired value to compare to.

You can define multiple criteria, which must all be satisfied in order for PingFederate to move a request to the next phase. A criterion is satisfied when the runtime value of the selected attribute matches or does not match the specified value, depending on the chosen comparison method. The multi-value contains ... or multi-value does not contain ... comparison methods are intended for attributes that can contain multiple values. Such a criterion is considered satisfied if one of the multiple values match or does not match the specified value. Values are compared verbatim. If you require complex evaluations, including conditional criteria or partial matching, define them using attribute mapping expressions.



Note:

All criteria defined must be satisfied, or evaluated as true, for a request to move forward, regardless of how the criteria were defined. As soon as one criterion fails, PingFederate rejects the request and returns an error message.

Steps

1. Go to Authentication # OAuth # IdP Adapter Grant Mapping and select your mapping. Click the Issuance Criteria tab.

2. From the **Source** list, select the attribute's source.

Depending on the selection, the Attribute Name list populates with associated attributes. See the following table for more information.

Source	Description
AccountLink	Select to evaluate the Local User ID value of the user.
	Displayed only if Account Linking is the selected identity mapping method. For more information, see .
Assertion	Select to evaluate attributes from the IdP connection.
Assertion or Provider Claims	Select to evaluate attributes from the IdP connection.
Context	Select to evaluate properties returned from the context of the transaction at runtime.
	Note:
	As the HTTP Request context value is retrieved as a Java object rather than text, attribute mapping expressions are more appropriate to evaluate and return values.
Extended Client Metadata	Select to evaluate OAuth client metadata.
JDBC, LDAP, or other types of datastore (if configured)	Select to evaluate attributes returned from a data source.
Mapped Attributes	Select to evaluate the mapped attributes.

- 3. From the Attribute Name list, select the attribute to be evaluated.
- 4. From the **Condition** list, select the comparison method.

Available methods:

- equal to
- equal to (case insensitive)
- equal to DN
- not equal to
- not equal to (case insensitive)
- not equal to DN
- multi-value contains
- multi-value contains (case insensitive)
- multi-value contains DN
- multi-value does not contain
- multi-value does not contain (case insensitive)
- multi-value does not contain DN



Note:

The first six conditions are intended for single-value attributes. Use one of the multi-value ... conditions for PingFederate to validate whether one of the attribute values matches the specified value. When an attribute has multiple values, using a single-value condition causes the criteria to fail. 5. In the **Value** field, enter the comparison value.



Values are compared verbatim. If you require complex evaluations, including conditional criteria or partial matching, define them using attribute mapping expressions. For more information, see Attribute mapping expressions on page 214.

6. In the **Error Result** field, enter a custom error message.

The value of this field is used by the error description protocol field. Using an error code in the Error Result field allows an application to process the code several ways, such as displaying an error message or e-mailing an administrator.

To use localized descriptions, enter a unique alias in the Error Result field, such as someIssuanceCriterionFailed. Insert the same alias with the desired localized text in the applicable language resource files, located in the <pf install>/pingfederate/server/ default/conf/language-packs directory.

If not defined, PingFederate returns ACCESS DENIED when the criterion fails at runtime.

- 7. Click Add.
- 8. Optional: Repeat to add more criteria.
- 9. If you require complex evaluations, including conditional criteria or partial matching, define them using attribute mapping expressions. For more information, see Attribute mapping expressions on page 214.
 - a. Click Show Advanced Criteria.
 - b. In the **Expression** field, enter the required expressions.
 - c. Optional: In the Error Result field, enter an error code or message.



Note:

If the expressions resolve to a string value instead of true or false, the returned value overrides the Error Result field value.

- d. Click Add.
- e. Optional: Click **Test**, enter values in the applicable fields, and verify the results.
- f. Optional: Repeat to add multiple criteria using attribute mapping expressions.

Reviewing the OAuth attribute mapping summary

On the **Summary** tab, review your configuration.

Steps

- To amend your configuration, click the corresponding tab title and then follow the steps to complete the task.
- To keep your changes, click Save.
- To discard your changes, click Cancel.

Managing authentication policy contract grant mapping

Use the Authentication Policy Contract Grant Mapping window to map values obtained from the authentication policy contract into the persistent grants. Persistent grants and any associated attributes and their values remain valid until the grants expire or until PingFederate explicitly revokes or cleans them up.

About this task

The USER_KEY attribute is the identifier of the persistent grants. The USER_NAME attribute presents the name shown to the resource owner on OAuth user-facing pages. If extended attributes are defined in System # OAuth Settings # Authorization Server Settings, configure a mapping for each attribute. You can optionally set up datastore queries to supplement values returned from the source. This mapping configuration is suitable for the Authorization Code and Implicit grant types.

Steps

Go to Authentication # OAuth # Authentication Policy Contract Grant Mapping and perform one of the following actions.

Action	Steps
Create a mapping	Select the source of the attributes from the list and click Add Mapping .
Modify an existing mapping	Select your mapping under Mappings .
Remove an existing mapping or cancel the removal request	Click Delete or Undelete under Action . Note: Before removing a mapping from your configuration, ensure that it is not used by your OAuth use cases. Any corresponding entries defined in Applications # OAuth # Access Token Mapping will also be removed.

Configuring policy contract attribute sources and user lookup

You can optionally set up datastore queries to supplement values returned from the policy contract attribute source.

Steps

- Go to Authentication # OAuth # Authentication Policy Grant Mapping and select your mapping, or click Add Mapping.
- On the Attribute Sources & User Lookup tab, perform one of the following actions. Choose from:
 - To set up datastore queries, click Add Attribute Source and follow the steps to complete the setup. For configuration steps, see Datastore query configuration on page 362.
 - To skip this configuration, click **Next**.

Fulfilling policy contract grant mapping

On the Contract Fulfillment tab, map authentication source values into persistent grants.

About this task

The USER_KEY attribute is the identifier of the persistent grants. The USER_NAME attribute presents the name shown to the resource owner on OAuth user-facing pages. If extended attributes are defined in System # OAuth Settings # Authorization Server Settings, configure a mapping for each attribute.



The USER KEY attribute values must be unique across all end users, because the USER KEY attribute is the user identifier to store and to retrieve persistent grants. For example, if you are configuring an OAuth attribute mapping on a SAML 2.0 identity provider (IdP) connection and the SAML SUBJECT attribute uniquely identifies all end users, you can map the SAML SUBJECT attribute to the USER KEY attribute.

Steps

1. On the Contract Fulfillment tab, select a source from the Source list, and then select or enter a value for each attribute in the contract.

Map each attribute from one of the following sources:

• Authentication Policy Contract

Populates the associated Value list with attributes associated with the APC.

Context

Values are returned from the context of the transaction at runtime.



If PERSISTENT GRANT LIFETIME is an extended attribute in System # OAuth Settings # Authorization Server Settings, you can set the lifetime of persistent grants based on the outcome of attribute mapping expressions, or the per-client **Persistent Grants Max Lifetime** setting.

- To set lifetime based on the per-client Persistent Grants Max Lifetime setting, select
 Context from the Source list and Default Persistent Grant Lifetime from the Value list.
- To set lifetime based on the outcome of attribute mapping expressions, select **Expression** as the source and enter an OGNL expression in the **Value** field.

If the expression returns a positive integer, the value represents the lifetime of the persistent grant in minutes.

If the expression returns the integer 0, PingFederate does not store the grant and does not issue a refresh token.

If the expression returns any other value, PingFederate sets the lifetime of the persistent grant based on the per-client **Persistent Grants Max Lifetime** setting.

 To set a static lifetime, select Text from the Source list and enter a static value in the Value field.

This is suitable for testing purposes, or cases where the persistent grant lifetime must always be set to a specific value.

As the **HTTP Request** context value is retrieved as a Java object rather than text, OGNL expressions are ideal to evaluate and return values.

• Extended Client Metadata

Values are returned from the client record.

LDAP/JDBC/Other (when a datastore is used)

Values are returned from your datastore. When you make this selection, the **Value** list populates with attributes from the datastore.

Expression (when enabled)

Provides more complex mapping capabilities, such as transforming incoming values into different formats. All of the variables available for text entries are also available for expressions.

• No Mapping

Ignores the Value field.

• Text

You can enter a text value only, or you can mix text with references to the unique user ID returned from the credentials validator, using the $\{attribute\}$ syntax. You can also enter values from your datastore, when applicable, using the $\{ds.attribute\}$ syntax, where attribute is any of the datastore attributes you have selected.

2. Click Next.

Defining issuance criteria for policy contract mapping

Individual attributes within policy contracts can further determine whether PingFederate approves or rejects requests. You can define those criteria to satisfy or you can choose to skip this configuration.

About this task

On the **Issuance Criteria** tab, define the criteria to satisfy in order for PingFederate to further process a request. Use this token authorization feature to conditionally approve or reject requests based on individual attributes.

Begin this optional configuration by choosing the source that contains the attribute to verify. Some sources are common to almost all use cases, such as **Mapped Attributes**. Other sources depend on the type of configuration, such as **JDBC**. Irrelevant sources are automatically hidden. Once you select a source, choose the attribute to verify. Depending on the selected source, the available attributes or properties vary. Specify the comparison condition and the desired value to compare to.



All criteria defined must be satisfied, or evaluated as true, for a request to move forward, regardless of how the criteria were defined. As soon as one criterion fails, PingFederate rejects the request and returns an error message.

Steps

1. From the **Source** list, select the attribute's source.

Depending on the selection, the Attribute Name list populates with associated attributes. See the following table for more information.

Source	Description
Authentication Policy Contract	Select to evaluate attributes from the authentication policy contract.
Context	Select to evaluate properties returned from the context of the transaction at runtime.
	Note: As the HTTP Request context value is retrieved as a Java object rather than text, attribute mapping expressions are more appropriate to evaluate and return values.
Extended Client Metadata	Select to evaluate OAuth client metadata.
JDBC, LDAP, or other types of datastore (if configured)	Select to evaluate attributes returned from a data source.
Mapped Attributes	Select to evaluate the mapped attributes.

2. From the **Attribute Name** list, select the attribute to be evaluated.

3. From the **Condition** list, select the comparison method.

Available methods:

- equal to
- equal to (case insensitive)
- equal to DN
- not equal to
- not equal to (case insensitive)
- not equal to DN
- multi-value contains
- multi-value contains (case insensitive)
- multi-value contains DN
- multi-value does not contain
- multi-value does not contain (case insensitive)
- multi-value does not contain DN



Note:

The first six conditions are intended for single-value attributes. Use one of the multi-value ... conditions for PingFederate to validate whether one of the attribute values matches the specified value. When an attribute has multiple values, using a single-value condition causes the criteria to fail.

4. In the **Value** field, enter the comparison value.



Note:

Values are compared verbatim. If you require complex evaluations, including conditional criteria or partial matching, define them using attribute mapping expressions. For more information, see Attribute mapping expressions on page 214.

5. In the **Error Result** field, enter a custom error message.

The value of this field is used by the error description protocol field. Using an error code in the Error Result field allows an application to process the code in several ways, such as displaying an error message or e-mailing an administrator.

To use localized descriptions, enter a unique alias in the Error Result field, such as someIssuanceCriterionFailed. Insert the same alias with the desired localized text in the applicable language resource files, located in the <pf install>/pingfederate/server/ default/conf/language-packs directory.

If not defined, PingFederate returns ACCESS DENIED when the criterion fails at runtime.

- 6. Click Add.
- 7. Optional: Repeat to add more criteria.

- 8. If you require complex evaluations, including conditional criteria or partial matching, define them using attribute mapping expressions. For more information, see Attribute mapping expressions on page 214.
 - a. Click Show Advanced Criteria.
 - b. In the **Expression** field, enter the required expressions.
 - c. Optional: In the **Error Result** field, enter an error code or message.



If the expressions resolve to a string value instead of true or false, the returned value overrides the Error Result field value.

- d. Click Add.
- e. Optional: Click **Test**, enter values in the applicable fields, and verify the results.
- f. Optional: Repeat to add multiple criteria using attribute mapping expressions.

Reviewing authentication policy contract mapping

On the **Summary** tab, review your configuration.

Steps

- To amend your configuration, click the corresponding tab title and then follow the steps to complete the task.
- To keep your changes, click Save.
- To discard your changes, click Cancel.

Managing resource owner credentials grant mapping

Use the Resource Owner Credentials Grant Mapping to map values obtained from the password credential validator instance into the persistent grants. Persistent grants and any associated attributes and their values remain valid until the grants expire or until PingFederate explicitly revokes or cleans them up.

About this task

The USER KEY attribute is the identifier of the persistent grants. If extended attributes are defined in System # OAuth Settings # Authorization Server Settings, configure a mapping for each attribute. You can optionally set up datastore queries to supplement values returned from the source. This mapping is intended for the Resource Owner Password Credential grant type.

Steps

Go to Authentication # OAuth # Resource Owner Credentials Grant Mapping and perform one of the following actions.

Action	Steps
Create a mapping	Select the source of the attributes from the list and click Add Mapping .
Modify an existing mapping	Select your mapping under Mappings .
Remove an existing mapping or cancel the removal request	Click Delete or Undelete under Action. Note:
	Before removing a mapping from your configuration, ensure that it is not used by your OAuth use cases. Any corresponding entries

Action	Steps
	defined in Applications # OAuth # Access Token Mapping will also be removed.

Configuring resource owner attribute sources and user lookup

You can optionally set up datastore queries to supplement values returned from the resource owner attribute source.

Steps

Go to Authentication # OAuth # Resource Owner Credentials Grant Mapping and select your mapping, or click Add Mapping.

Choose from:

- To set up datastore queries, click Add Attribute Source on the Attribute Sources and User Lookup tab. For configuration steps, see *Datastore query configuration* on page 362.
- To skip this optional configuration, click **Next**.

Fulfilling resource owner credentials grant mapping

On the **Contract Fulfillment** tab, map authentication source values into persistent grants.

About this task

The USER KEY attribute is the identifier of the persistent grants. If extended attributes are defined in System # OAuth Settings # Authorization Server Settings, configure a mapping for each attribute.



Important:

The USER KEY attribute values must be unique across all end users, because the USER KEY attribute is the user identifier to store and to retrieve persistent grants. For example, the samaccountName attribute value of an end user in one domain might match that of another end user in another domain. In this case, you can map the Subject DN attribute to the USER KEY attribute.

Steps

1. On the Contract Fulfillment tab, select a source from the Source list, and then select or enter a value for each attribute in the contract.

Map each attribute from one of the following sources:

• Password Credential Validator

When selected, the associated Value list populates with attributes associated with the credentialvalidation instance.

Context

Values are returned from the context of the transaction at runtime.



Note:

If PERSISTENT GRANT LIFETIME is an extended attribute in System # OAuth Settings # Authorization Server Settings, you can set the lifetime of persistent grants based on the outcome of attribute mapping expressions, or the per-client **Persistent Grants Max Lifetime** setting.

- To set lifetime based on the per-client Persistent Grants Max Lifetime setting, select
 Context from the Source list and Default Persistent Grant Lifetime from the Value list.
- To set lifetime based on the outcome of attribute mapping expressions, select **Expression** as the source and enter an OGNL expression in the **Value** field.

If the expression returns a positive integer, the value represents the lifetime of the persistent grant in minutes.

If the expression returns the integer 0, PingFederate does not store the grant and does not issue a refresh token.

If the expression returns any other value, PingFederate sets the lifetime of the persistent grant based on the per-client **Persistent Grants Max Lifetime** setting.

 To set a static lifetime, select Text from the Source list and enter a static value in the Value field.

This is suitable for testing purposes, or cases where the persistent grant lifetime must always be set to a specific value.

As the **HTTP Request** context value is retrieved as a Java object rather than text, OGNL expressions are ideal to evaluate and return values.

• Extended Client Metadata

Values are returned from the client record.

LDAP/JDBC/Other (when a datastore is used)

Values are returned from your datastore. When selected, the **Value** list is populates with attributes from the datastore.

Expression (when enabled)

Provides more complex mapping capabilities, such as transforming incoming values into different formats. All of the variables available for text entries are also available for expressions.

• No Mapping

Ignores the Value field.

Text

You can enter a text value only, or you can mix text with references to the unique user ID returned from the credentials validator, using the $\{attribute\}$ syntax. You can also enter values from your datastore, when applicable, using the $\{ds.attribute\}$ syntax, where attribute is any of the datastore attributes you have selected.

2. Click Next.

Defining issuance criteria for resource-owner credentials mapping

Individual attributes within policy contracts can further determine whether PingFederate approves or rejects requests. You can define those criteria to satisfy or you can choose to skip this configuration.

About this task

On the **Issuance Criteria** tab, define the criteria to satisfy in order for PingFederate to further process a request. Use this token authorization feature to conditionally approve or reject requests based on individual attributes.

Begin this optional configuration by choosing the source that contains the attribute to verify. Some sources are common to almost all use cases, such as **Mapped Attributes**. Other sources depend on the type of configuration, such as **JDBC**. Irrelevant sources are automatically hidden. Once you select a source, choose the attribute to verify. Depending on the selected source, the available attributes or properties vary. Specify the comparison condition and the desired value to compare to.



Note:

All criteria defined must be satisfied, or evaluated as true, for a request to move forward, regardless of how the criteria were defined. As soon as one criterion fails, PingFederate rejects the request and returns an error message.

Steps

- 1. Go to Authentication # OAuth # Resource Owner Credentials Grant Mapping and select your mapping, or click Add Mapping.
- 2. On the Issuance Criteria tab, select the attribute's source from the Source list.

criteria or partial matching, define them using attribute mapping expressions.

Depending on the selection, the Attribute Name list populates with associated attributes. See the following table for more information.

Source	Description
Context	Select to evaluate properties returned from the context of the transaction at runtime.
	Note:
	As the HTTP Request context value is retrieved as a Java object rather than text, attribute mapping expressions are more appropriate to evaluate and return values.
Extended Client Metadata	Select to evaluate OAuth client metadata.
JDBC, LDAP, or other types of datastore (if configured)	Select to evaluate attributes returned from a data source.
Mapped Attributes	Select to evaluate the mapped attributes.
Password Credential Validator	Select to evaluate attributes from the Password Credential Validator instance.

3. From the **Attribute Name** list, select the attribute to be evaluated.

4. From the **Condition** list, select the comparison method.

Available methods:

- equal to
- equal to (case insensitive)
- equal to DN
- not equal to
- not equal to (case insensitive)
- not equal to DN
- multi-value contains
- multi-value contains (case insensitive)
- multi-value contains DN
- multi-value does not contain
- multi-value does not contain (case insensitive)
- multi-value does not contain DN



Note:

The first six conditions are intended for single-value attributes. Use one of the multi-value ... conditions for PingFederate to validate whether one of the attribute values matches the specified value. When an attribute has multiple values, using a single-value condition causes the criteria to fail.

5. In the **Value** field, enter the comparison value.



Note:

Values are compared verbatim. If you require complex evaluations, including conditional criteria or partial matching, define them using attribute mapping expressions. For more information, see Attribute mapping expressions on page 214.

6. In the **Error Result** field, enter a custom error message.

The value of this field is used by the error description protocol field. Using an error code in the Error Result field allows an application to process the code in several ways, such as displaying an error message or e-mailing an administrator.

To use localized descriptions, enter a unique alias in the Error Result field, such as someIssuanceCriterionFailed. Insert the same alias with the desired localized text in the applicable language resource files, located in the <pf install>/pingfederate/server/ default/conf/language-packs directory.

If not defined, PingFederate returns ACCESS DENIED when the criterion fails at runtime.

- 7. Click Add.
- 8. Optional: Repeat to add more criteria.

- 9. If you require complex evaluations, including conditional criteria or partial matching, define them using attribute mapping expressions. For more information, see Attribute mapping expressions on page 214.
 - a. Click Show Advanced Criteria.
 - b. In the **Expression** field, enter the required expressions.
 - c. Optional: In the **Error Result** field, enter an error code or message.



Note:

If the expressions resolve to a string value instead of true or false, the returned value overrides the Error Result field value.

- d. Click Add.
- e. Optional: Click **Test**, enter values in the applicable fields, and verify the results.
- f. Optional: Repeat to add multiple criteria using attribute mapping expressions.

Reviewing the resource owner credentials mapping

On the **Summary** tab, review your configuration.

Steps

- To amend your configuration, click the corresponding tab title and then follow the steps to complete the task.
- To keep your changes, click Save.
- To discard your changes, click Cancel.

Token mapping

You can configure how PingFederate maps attributes to OAuth access tokens and OpenID Connect ID tokens.

To configure how PingFederate maps attributes to OAuth access tokens, go to Applications # OAuth # Access Token Mappings. You can map attributes to access tokens from various context sources, depending on how users or clients authenticate. This is the second stage of mapping OAuth attributes to access tokens. The first stage is mapping OAuth attributes to persistent grant contracts. For more information, see Mapping OAuth attributes on page 81 and Managing access token mappings on page 561.

OpenID Connect policies include a map of attributes to ID tokens. To configure an OpenID Connect policy, go to Applications # OAuth # OpenID Connect Policy Management. For more information, see Configuring OpenID Connect policies on page 572.

Access token management

PingFederate supports multiple access token management (ATM) instances. You can configure different access token policies and attribute contracts for different OAuth clients. You can also control validation of access tokens to one or more resource servers.

When defining an ATM instance, you can customize various settings, including token format, lifetime, session validation settings, and attribute contract for this instance. You can also limit the ATM instance to a list of resource URIs, a set of clients in an access control list (ACL), or both.

For example, you can use the ACL to limit which clients can obtain access tokens from a particular ATM instance. You can also add a resource server client to the ACL of multiple ATMs instances, so that only the resource server client can submit token validation requests for access tokens issued by those ATM instances.

When there are multiple ATM instances, OAuth clients can specify the desired ATM instance by providing the access token manager id ATM ID or an aud resource URI in their requests to the PingFederate OAuth authorization server at the /as/authorization.oauth2 authorization endpoint, the /as/ token.oauth2 token endpoint, and the /as/introspect.oauth2 introspection endpoint. For resource server clients, you can configure on a per-client basis whether a resource server client must specify the desired ATM instance in its token validation requests at runtime. For more information, see *Configuring an OAuth client* on page 515.

At runtime, the PingFederate OAuth authorization server uses the following rules to determine which ATM instances to use:

- PingFederate limits the eligible ATM instances to those that are available in the context of the request.
 For most requests, these are instances that have an attribute mapping defined in the Access Token
 Mapping window. For OAuth Assertion Grant requests, it is the set of instances for which a mapping is
 defined in the IdP connection. If configured, the ACL can also limit which ATM instances are eligible.
- 2. If the request comes with an access_token_manager_id or aud parameter, PingFederate uses the information to determine the applicable ATM instance.
- 3. If the request does not come with either parameter, for OAuth clients supporting the OpenID Connect protocol by including the openid scope value, PingFederate uses the ATM instance specified by the OpenID Connect policy associated with the client. For resource server clients, you can optionally configure PingFederate to use any eligible ATM instances for the purpose of token validation.
- 4. If the request comes with neither of the two parameters nor the openid scope, PingFederate uses the default ATM instance of the client if configured, or the default ATM instance defined for the installation if eligible. For token validation requests, if resource server clients do not provide either the access_token_manager_id or aud parameter in their requests and the resource server clients have not been configured to validate against any eligible ATM instances, the same logic applies.

If no match can be found in the eligible list of ATMs, PingFederate aborts the request.

Managing access token management instances

Use the **Access Token Management** window to specify how the PingFederate OAuth AS manages access tokens.

Steps

- 1. Go to Applications # OAuth # Access Token Management.
- 2. In the Access Token Management window, choose from the following options.

Option	Description
Create New Instance	Configure a new instance
<existing instance="" link=""> under Instance Name</existing>	Modify an existing instance
Check Usage	Review the usage of an existing instance
Delete or Undelete	Remove an existing instance or cancel the removal

Defining an access token management instance

Define your access token management instance in the **Type** tab. This capability allows you to configure different access token policies and attribute contracts for different OAuth clients. It also provides a means to control validation of access tokens to one or more resource servers.

Steps

- 1. Go to Applications # OAuth # Access Token Management and click Create New Instance.
- 2. On the **Type** tab, enter a name in the **Instance Name** field and an ID in the **Instance ID** field.
- 3. From the **Type** list, select the plugin type of the access token management instance.

The **Type** list varies depending on the plugins deployed on your server. For information about adding a customized plugin, contact the *Ping Identity Support Center*.

4. Optional: Select a Parent Instance from the list.

Use this option when creating an instance that is similar to an existing one. The child instance inherits the configuration of its parent. You can also override one or more settings during the setup. Select the **Override ...** check box and make the adjustments as needed in one or more subsequent windows.

Configuring an access token management instance

This configuration varies depending on the type of the access token manager.

Steps

See subsequent topics for configuration steps.

Configuring reference token management

Access tokens that use the reference token data model provide a reference to a set of attributes. The resource server must de-reference the access tokens for the corresponding identity and security information at the OAuth authorization server that issued them. PingFederate is the authorization server.

About this task

The reference token data model supports both adaptive clustering and directed clustering. For adaptive clustering, PingFederate shares token information across a replica set. If region identifiers are defined, PingFederate shares token information across replica sets in multiple regions. You can optionally override this default behavior in the configuration file for adaptive clustering. For directed clustering, PingFederate shares token information among all engine nodes, despite any state server or subcluster setup.

Steps

- 1. Go to Applications # OAuth # Access Token Management and click Create New Instance.
- 2. On the **Instance Configuration** window, modify the default values as needed.

The following table describes each field.

Field	Description
Token Length (Required)	The number of characters that PingFederate uses to define the token reference. Increasing the length enhances token security.
(rioquiiou)	The default value is 28 characters. The minimum and maximum values are 22 and 256, respectively.
Token Lifetime	The amount of time in minutes that an access token is considered valid.
(Required)	The default value is 120 minutes.
Lifetime Extension Policy	Indicates whether PingFederate should reset the lifetime of an access token each time the token is validated, subject to the values defined in the Maximum Token Lifetime and Lifetime Extension Threshold Percentage fields.
	The options are:
	 No Extension Tokens Not Backed by Persistent Access Grants (Transient Grants) All Tokens
	The default selection is No Extension .

Field	Description
Maximum Token Lifetime	Defines an absolute maximum token lifetime for use with the Lifetime Extension Policy setting, in minutes. When configured, the lifetime of access tokens can be extended but not beyond the configured value. Any value, if specified, must be greater than or equal to the value specified in the Token Lifetime field.
	This optional field has no default value.
Lifetime Extension Threshold Percentage (Required)	When PingFederate is deployed in a cluster and token-lifetime extension is enabled, there must be a cluster-group remote procedure call (RPC) to extend the life of a token.
(Kequireu)	To limit RPC overhead, this setting suspends the calls until the remaining time is less than the chosen value, as a percentage of token lifetime. For example, if the token lifetime is 60 minutes and the Lifetime Extension Threshold Percentage value is 30 percent, the lifetime will not be extended until the remaining time is less than 18 minutes. This option can drastically reduce RPC traffic between nodes, while still supporting a lifetime extension policy.
	The default value is 30 percent.
Advanced Fields	
Mode for Synchronous RPC	Synchronous RPC calls occur when a node receives a verification request for a token it does not recognize, and for token issuance.
	When Majority of Nodes is selected, the server waits for the majority of recipients to respond. It also eliminates the need for a complete state synchronization at startup.
	When All Nodes is selected, the server waits for all recipients to respond.
	The default selection is Majority of Nodes .
RPC Timeout (Required)	The timeout value between cluster nodes during synchronous communication, in milliseconds. The recommended value ranges from 100 milliseconds to 1000, or 1 second.
	The default value is 500 milliseconds.
Expand Scope Groups	Determines whether to expand scope groups into their corresponding scopes in the access token contents and introspection response.
	This check box is not selected by default.

Configuring JSON token management

JSON web token (JWT) bearer access tokens are secure and self-contained tokens. This allows the target resource server to validate the access tokens locally or to send the access tokens to PingFederate for validation.

About this task

This configuration uses either symmetric keys or asymmetric signing-certificate keys for token security. To facilitate rollover of keys when they expire, multiple entries are allowed for either signing mechanism. This token data model is suitable for both standalone and clustered environments.

Steps

1. Go to Applications # OAuth # Access Token Management and click Create New Instance.

- 2. On the **Instance Configuration** tab, add one or more symmetric keys, signing certificates, or both.
 - a. Click Add a new row to..., or click Update to modify an existing entry.



The **Key ID** field values must be unique across all JSON token management instances, including child instances.

b. If you have not yet created or imported your certificate into PingFederate, click **Manage Signing Certificates** and complete the task.



Note:

To use an RSA-based algorithm for JSON web signature (JWS), the key size of the signing certificate must be at least 2,048 bits. For an EC-based JWS algorithm, the key size depends on the chosen algorithm.

3. Change or select entries as needed.

For more information about JSON web algorithms (JWAs), see the *JSON Web Algorithms* specification.

Field	Description
Token Lifetime	The amount of time that an access token is considered valid, in minutes.
(Required)	The default value is 120 minutes.
Use Centralized Signing Key	Select this option to use a centralized key when signing JWTs using an RSA-based or EC-based algorithm.
	When this option is selected and static keys are not enabled, PingFederate manages and rotates a set of keys, and uses the current key corresponding to the selected signing algorithm to sign the JWT.
	When this option is selected and static keys are enabled, PingFederate uses the current key corresponding to the selected signing algorithm to sign the JWT.
	When this option is selected, an OAuth client that has been configured to use JWT access tokens through this ATM instance can retrieve the key it needs to validate the digital signature by contacting PingFederate at the $/\mathrm{pf/JWKS}$ endpoint.
JSON Web Signature (JWS) configuration
JWS Algorithm	The hash-based message authentication code (HMAC) or the signing algorithm (EC or RSA) used to protect the integrity of the token. If PingFederate is deployed to run in a Java 11 runtime environment or integrated with a hardware security module (HSM), additional RSASSA-PSS signing algorithms become available for selection. For more information on HSM integration, see <i>Supported hardware security modules</i> on page 169.
	Required if an asymmetric algorithm is selected in the JWE Algorithm list.

Field	Description
Active Symmetric Key ID	The ID of the symmetric key to use when producing JWTs using an HMAC-based algorithm.
	Required if an HMAC-based JWS algorithm is selected in the JWS Algorithm list.
Active Signing Certificate Key ID	The ID of the key pair and certificate to use when producing JWTs using an EC-based or RSA-based algorithm.
	Required if an EC-based or RSA-based JWS algorithm is selected in the JWS Algorithm list.
JSON Web Encryption	(JWE) configuration
JWE Algorithm	The algorithm used to encrypt or otherwise determine the value of the content encryption key.
	PingFederate supports symmetric algorithms, such as Direct Encryption with symmetric key, AES Key Wrap, and AES-GCM key encryption, and asymmetric algorithms, such as ECDH-ES, ECDH-ES Key Wrap, and RSAES OAEP.
JWE Content Encryption Algorithm	The content encryption algorithm used to perform authenticated encryption on the plain text payload of the token.
	Required if an algorithm is selected in the JWE Algorithm list.
Active Symmetric	The ID of the key to use when using a symmetric encryption algorithm.
Encryption Key ID	Required if a symmetric algorithm is selected in the JWE Algorithm list.
Asymmetric Encryption Key	An asymmetric encryption public key from your partner, which can be in either JSON web key (JWK) format or a certificate.
	Applicable only if an asymmetric algorithm is selected from the JWE Algorithm list.
	Note:
	You can only specify an asymmetric encryption key here or the partner's JSON web key set (JWKS) endpoint in the Asymmetric Encryption JWKS URL field.
Asymmetric Encryption JWKS	The HTTPS URL of a JWKS endpoint that provides a list of one or more public keys for encryption.
URL	Applicable only if an asymmetric algorithm is selected from the JWE Algorithm list.
	Note:
	You can only specify an asymmetric encryption JWK URL here or the asymmetric encryption public key from your partner in the Asymmetric Encryption Key field.
Advanced fields	

Field	Description
Include Key ID Header	When selected, the key ID is used in the kid header parameter for the
Parameter	token.
	This check box is selected by default.
Include X.509 Thumbprint Header	When selected, the X.509 certificate thumbprint is used in the ${\tt x5t}$ header parameter for the token.
Parameter	This check box is not selected by default.
Default JWKS URL Cache Duration	When an asymmetric encryption JWKS URL is specified, if the remote server does not contain any cache directives in its response, PingFederate only caches the content for 720 minutes (12 hours).
	Note:
	When this threshold is reached or if the cache directives indicate that the content has expired at runtime, PingFederate contacts the remote server to refresh the list of encryption keys from the partner.
Include JWE Key ID header parameter	When selected, indicates whether the key ID (kid) header parameter will be included in the encryption header of the token, which can help identify the appropriate key during decryption.
	This check box is selected by default.
Include JWE X.509 Thumbprint Header Parameter	When selected, the X.509 certificate thumbprint is used as the x5t header parameter value in the encryption header of the token. This can help identify the appropriate key during decryption.
	This check box is not selected by default.
Client ID Claim Name	The name of a JWT claim used to represent the OAuth client ID.
	The default value is client_id.
	If the field value is empty, PingFederate will not include the client ID of the requesting client in the self-contained tokens. If clients may use the UserInfo endpoint to retrieve additional claims about the users, see <i>UserInfo endpoint</i> on page 1085 for more information.
Scope Claim Name	The name of a JWT claim used to represent the scope of the grant.
	The default value is scope.
	If the field value is empty, PingFederate will not include any scope information in the self-contained token. If clients may use the UserInfo endpoint to retrieve additional claims about the users, see <i>UserInfo endpoint</i> on page 1085 for more information.
Space Delimit Scope Values	When selected, indicates scope strings will be delimited by spaces rather than represented as a JSON array.
	This check box is not selected by default.

Field	Description
Issuer Claim Value	The value of the Issuer claim (iss) in the JWT. If left blank, this field is omitted.
	Additionally, you may extend the contract of the access tokens with an attribute named iss on the Access Token Attribute Contract tab. When mapping attribute values from authentication sources to the access tokens issued by this ATM instance, the value you specify on the Contract Fulfillment tab overrides the value here.
Audience Claim Value	The value of the Audience claim (aud) in the JWT. If left blank, this field is omitted.
	When no value is specified, PingFederate does not validate the aud value, if any is included in the access token.
	You can also extend the contract of the access tokens with an attribute named aud in the Access Token Attribute Contract tab. When mapping attribute values from authentication sources to the access tokens issued by this ATM instance, the value specified in the Contract Fulfillment window overrides this value.
JWT ID Claim Length	Indicates the number of characters of the JWT ID (jti) claim in the JWT.
	The default value is 0, meaning no claim is included.
Access Grant GUID Claim Name	The name of the JWT claim used to carry the persistent access grant GUID. If left blank, this field is omitted.
	If the claim is present during validation, PingFederate checks the grant database to ensure the grant is still valid.
	Note:
	This use case requires that the RS must send the JWT bearer access tokens to PingFederate for validation.
JWKS Endpoint Path	The path on the PingFederate server to publish a JWKS with the keys and certificates that the partners can use for signature verification. Optional when an algorithm is selected in the JWS Algorithm list. If specified, the path must begin with a forward slash, such as /oauth/jwks.
	The resulting URL is https:// <i><pf_host></pf_host></i> : <i><pf.https.port></pf.https.port></i> /ext/ <i><jwks endpoint="" path=""></jwks></i> .
	The path must be unique across all plugin instances, including any child instances.
JWKS Endpoint Cache Duration	Informs the clients of the duration that they could cache the content from the JWKS endpoint path. Applicable only if the JWKS Endpoint Path field is configured.
	The default is 720 minutes, or 12 hours.
Publish Key ID X.509 URL	Indicates whether certificates will be made accessible by the key ID at https:// <pf_host>:<pf.https.port>/ext/oauth/x509/kid?v=<id>.</id></pf.https.port></pf_host>
	This check box is not selected by default.

Field	Description
Publish Thumbprint X.509 URL	Indicates whether certificates will be made accessible by thumbprint at https:// <pf_host>:<pf.https.port>/ext/oauth/x509/x5t?v=<base64url encoded="" sha-1="" thumbprint="">.</base64url></pf.https.port></pf_host>
	This check box is not selected by default.
Expand Scope Groups	Determines whether to expand scope groups into their corresponding scopes in the access token contents and introspection response.
	This check box is not selected by default.
Type Header Value	Indicates the value of the Type (typ) header in the JWT. If you do not specify a header, it is omitted.

Managing session validation settings

When an OAuth client presents an access token for validation, PingFederate acts as an OAuth authorization server and checks the expiration and the other aspects of the access token. If the validation fails, PingFederate returns an invalid grant error to the client.

Before you begin

The session validation features require authentication sessions. You must enable authentication sessions for either all authentication sources or the authentication source associated with the OAuth use cases.

About this task

When PingFederate authentication sessions are enabled, you can optionally configure the access token validation process to evaluate the authentication sessions of the users, or resource owners, before returning the validation results to the clients. Depending on the features selected on the Session Validation tab, PingFederate might return an invalid grant error if the associated authentication session has timed out, expired, is not found, or has been revoked. You can also configure PingFederate to extend the authentication sessions upon successful validations.

When any session validation features are enabled, the associated session identifier (pi.sri) becomes available through the access tokens. For reference-style access tokens, PingFederate returns the associated session identifier in the response if the access token is valid. For JSON web token (JWT)based access tokens, the session identifier is part of the access token. Through the session identifier, an OAuth client can contact the Session Management API and Session Revocation API endpoints to query the status of an authentication session, or to extend or revoke an authentication session.

The session validation features let you combine the status of access tokens and user authentication sessions. Because you can independently enable each feature per access token management (ATM) instance, you can customize unique API and web single sign-on (SSO) behaviors for your OAuth clients and users.



Important:

The session validation features are most suitable for clients using the implicit grant type, which does not use refresh tokens. Clients using the authorization code grant type can still use session validation. However, they can only refresh their access tokens through ATM instances that have the session validation features disabled, other than the Include session identifier in access token feature.

When a client using the authorization grant type has an access token and sends a refresh token to an ATM instance, the following rules apply:

• If the ATM instance has Include session identifier in access token enabled and the other session validation features disabled, then the ATM instance can issue a new access token. If the original access token contains a session identifier, the new token will contain the same session identifier.

- If the ATM instance has any session validation feature enabled other than the Include session identifier in access token feature, then PingFederate returns an unsupported grant type error.
- If the ATM instance has Include session identifier in access token disabled and the other session validation features disabled, then the ATM instance can issue a new access token, but with no session identifier.

Steps

- 1. Go to Applications # OAuth # Access Token Management.
- 2. Select the applicable ATM instance or click **Create New Instance**.

If you are creating a new ATM instance, complete the required fields in the Type and Instance Configuration tabs.

3. On the **Session Validation** tab, select the check box for each relevant feature.



Important:

If authentication sessions are not enabled, you can still select features on this tab, but access token validation might fail.

If this is a child ATM instance, select the Override Session Validation Settings check box and edit as needed.

Each feature is independent of each other. The following table describes each feature.

Feature	Description
Include session identifier in access token	When selected, the ATM instance includes the value of the pi.sri session identifier in the access tokens it issues. An OAuth client that is allowed to access the session management API can get information about sessions associated with the session identifier. The client can also request PingFederate to extend or revoke the sessions. For more information, see Session Management API on page 1120.

Feature	Description
Check for valid authentication session	When selected, an access token is considered invalid unless the user has a valid authentication session. If the user does not have a valid session, PingFederate returns an invalid_grant error.
	An authentication session is invalid when one of the following conditions applies:
	 The authentication session has timed out based on the Idle Timeout field value in Authentication # Policies # Sessions.
	 The authentication session has expired based on the Max Timeout field value in Authentication # Policies # Sessions.
	 The authentication session is not found, such as if the user has logged out.
	You can also optimize the access token lifetime.
	Note:
	If this ATM instance issues internally managed reference tokens, match the Token Lifetime value in the Instance Configuration tab to the Idle Timeout value. If you specify a Maximum Token Lifetime value on the Instance Configuration tab, ensure that the value matches that of the Max Timeout field.
	If this ATM instance issues JWT-based access tokens, match value of the Token Lifetime field to that of the Max Timeout field.
Check session revocation status	When selected, PingFederate verifies whether the session identifier has been added to the revocation list. If the session has been revoked, PingFederate returns an invalid_grant error.
	An authentication session can be revoked through the front-channel or the back-channel.
Update authentication session activity	When selected, if the access token is valid, PingFederate also extends the lifetime of the authentication session by the Idle Timeout field value in Authentication # Policies # Sessions .
	For externally stored authentication sessions, this operation only sends updates to the external storage when the remaining idle timeout window is less than 75%.

Defining the access token attribute contract

On the **Access Token Attribute Contract** tab, define the attribute contract for the access tokens issued by this access token management (ATM) instance.

About this task

You must enter at least one attribute. For auditing purposes, an attribute can be chosen as the subject.

Steps

1. Go to Applications # OAuth # Access Token Management and select your instance, or click Create New Instance.

2. On the Access Token Attribute Contract tab, use the Extend the Contract field and the Add button to add one or more attributes.

For JSON web token (JWT) bearer access tokens, you can extend the attribute contract with the following attributes.

Attribute	Description
iss	Adds the Issuer claim (iss) to the access token.
	When mapping attribute values from authentication sources to the access tokens issued by this ATM instance, the value specified on the Access Token Attribute Contract tab overrides any Issuer Claim Value defined on the Instance Configuration tab.
aud	Adds the Audience claim (aud) to the access token.
	When mapping attribute values from authentication sources to the access tokens issued by this ATM instance, the value you specify on the Access Token Attribute Contract tab overrides any Audience Claim Value defined on the Instance Configuration tab.
ехр	Extends the value of the Expire claim (exp) by the specified value in seconds.
	Note:
	Define the Expire claim with the Token Lifetime setting in the Instance Configuration tab.
The Client ID Claim Name field value, the Scope Claim Name field value, or the Access Grant GUID Claim Name field value defined on the Instance Configuration tab of this ATM instance.	When mapping attribute values from authentication sources to the access tokens issued by this ATM instance, the values defined in the Access Token Attribute Contract tab override the value of the client ID, the scope, or the persistent access grant GUID.

3. Select an attribute from the Subject Attribute Name list.

Result:

When recording OAuth transactions in the audit log, populates the subject field with values from this attribute specifically for token introspection and token validation using the <code>validate_bearer</code> grant type.

Managing resource URIs

When sending its request to the authorization endpoint on the PingFederate OAuth authorization server, an OAuth client can optionally include the requested resource in the aud query parameter.

About this task

If the client is sending a token exchange request, it can specify an access token manager in the **resource** query parameter.

You can specify a list of resource URIs that PingFederate OAuth authorization server can use to select this access token management instance when the aud or resource query parameter is provided.



The resource URIs must correspond to the resource that the resource server expects.

Steps

- Go to Applications # OAuth # Access Token Management and select your instance, or click Create New Instance.
- 2. On the **Resource URI** tab, perform one of the following actions.

Action	Steps
Add a new entry	Enter the desired value in the Resource URIs field, and then click Add .
Modify an existing entry	Use the Edit, Update, and Cancel buttons.
Remove an existing entry	Use the Delete and Undelete buttons.

Defining access control

On the **Access Control** tab, you can restrict which OAuth clients are allowed to use this access token management instance.

Steps

- Go to Applications # OAuth # Access Token Management and select your ATM instance or click Create New Instance.
- 2. On the Access Control tab, select the Restrict Allowed Clients check box.
- 3. Select a client from the Allowed Clients list, and then click Add.

Repeat this step to select additional clients as needed.

Result

To remove a client from the **Allowed Clients** list or to cancel the removal request, click **Delete** or **Undelete** under **Action**.

To disable access control by clients altogether, clear the Restrict Allowed Clients check box.

Reviewing the access token management configuration

On the **Summary** tab, review your access token management configuration.

Steps

- To amend your configuration, click the corresponding tab title and then follow the steps to complete
 the task.
- To keep your changes, click Save.
- To discard your changes, click Cancel.

Managing access token mappings

In this required configuration, map attributes to be requested from the OAuth resource server into the access token, the token attribute contract.

About this task

When mapping a default context, define how PingFederate maps values into the attributes based on the persistent-grant user_key, and any extended attributes defined in System # OAuth Settings # Authorization Server Settings. PingFederate acts as an OAuth authorization server.

When a specific context is selected, you can map attributes from the selected context, specifically the chosen IdP adapter instance, Password Credential Validator instance, or authentication policy contract, into the access tokens. You can also map attributes from an IdP connection with an OAuth attribute

mapping configuration or an authentication policy contract mapping configuration. You can configure a mapping for clients using the client credential grant type.

The mapping used at runtime depends on the authentication context of the original grant. If the authentication context results in a match, PingFederate uses that specific mapping. Otherwise, it uses the default mapping for the applicable access token manager instance.



The Access Token Mapping window becomes available after at least one access token manager (ATM) instance has been configured in Applications # OAuth # Access Token Management.

Steps

Go to Applications # OAuth # Access Token Management.

Action	Steps
Create a mapping	Select the source of the attributes from the Context list and the target ATM instance from the Access Token Manager list, and then click Add Mapping .
Modify an existing mapping	Select it by its name under Mappings .
Remove an existing mapping or to cancel the removal request	Click Delete or Undelete under Action .
removal request	Note:
	Before removing an existing mapping from your configuration, ensure that it is not used by your OAuth use cases.

Configuring access token attribute sources and user lookup

You can optionally set up datastore queries to supplement values returned from the access token attribute source.

Steps

Go to Applications # OAuth # Access Token Mapping and select your mapping, or click Add Mapping. Choose from:

- To set up datastore queries, click Add Attribute Source and complete the task in the Attribute Sources & User Lookup tab. For configuration steps, see Datastore query configuration on page
- To skip this optional configuration, click Next.

Configuring access token fulfillment

On the Contract Fulfillment tab, map values into the token attribute contract to be included or referenced in the access token.

Steps

1. Choose a source from the Source list, and then select a value from the Value list for each attribute in the contract, or enter your own.

Map each attribute from one of the following sources:

- Client Credentials, IdP Adapter, IdP Connection, Password Credential Validator, or Token **Exchange Processor Policy**

Depending on the selections under Context in the Access Token Attribute Mapping tab. you can map attributes from that specific authentication system. Select the corresponding context under **Source** and the desired attribute under **Value**.

Persistent Grant

When selected, the associated Value list is populated with the USER KEY and extended attributes from the persistent access-token grant.

Context

Values are returned from the context of the transaction at runtime.



Note:

The **HTTP Request** context value is retrieved as a Java object rather than text. For this reason, OGNL expressions are preferred to evaluate and return values.

Select **Expression** under **Source**, and then click **Edit** to enter an expression.

The HTTP RequestJava object retrieves the authentication method that a client uses, or the private key JWT for client authentication if the client uses the private key jwt authentication method.. For sample expressions, see Expressions for OAuth and OpenID Connect uses cases on page 219.

If the **Expression** selection is not available, you can enable it by editing the org.sourceid.common.ExpressionManager.xml file in the <pf install>/ pingfederate/server/default/data/config-store directory.

Extended Client Metadata

Values are returned from the client record.

LDAP/JDBC/Other

Values are returned from your datastore, if used. When you make this selection, the Value list populates with attributes from the datastore.

Expression

When enabled, this option provides more complex mapping capabilities, such as transforming incoming values into different formats. All of the variables available for text entries are also available for expressions.

No Mapping

This option ignores the **Value** field, causing no value selection to be necessary.

Text

The value is what you enter. This can be text only, or you can mix text with references to the **USER KEY using the \${USER KEY} syntax.**

When applicable, you can also enter values from your datastore using the \${ds.attribute} syntax, where attribute is any of the datastore attributes you have selected.

Defining issuance criteria for access token mapping

Individual attributes within policy contracts can further determine whether PingFederate approves or rejects requests. You can define those criteria to satisfy or you can choose to skip this configuration..

About this task

On the Issuance Criteria tab, define the criteria to satisfy in order for PingFederate to further process a request. Use this token authorization feature to conditionally approve or reject requests based on individual attributes.

Begin this optional configuration by choosing the source that contains the attribute to verify. Some sources are common to almost all use cases, such as Mapped Attributes. Other sources depend on the type of configuration, such as JDBC. Irrelevant sources are automatically hidden. Once you select a source, choose the attribute to verify. Depending on the selected source, the available attributes or properties vary. Specify the comparison condition and the desired value to compare to.

You can define multiple criteria, which must all be satisfied in order for PingFederate to move a request to the next phase. A criterion is satisfied when the runtime value of the selected attribute matches or does not match the specified value, depending on the chosen comparison method. The multi-value contains ... or multi-value does not contain ... comparison methods are intended for attributes that can contain multiple values. Such a criterion is considered satisfied if one of the multiple values match or does not match the specified value. Values are compared verbatim. If you require complex evaluations, including conditional criteria or partial matching, define them using attribute mapping expressions.



Note:

All criteria defined must be satisfied, or evaluated as true, for a request to move forward, regardless of how the criteria were defined. As soon as one criterion fails, PingFederate rejects the request and returns an error message.

Steps

- 1. Go to Applications # OAuth # Access Token Mapping and select your mapping, or click Add Mapping.
- 2. On the Issuance Criteria tab, select the attribute's source from the Source list.

Depending on the selection, the **Attribute Name** list populates with associated attributes. See the following table for more information.

Source	Description
Context	Select to evaluate properties returned from the context of the transaction at runtime.
	Note:
	As the HTTP Request context value is retrieved as a Java object rather than text, attribute mapping expressions are more appropriate to evaluate and return values.
Extended Client Metadata	Select to evaluate OAuth client metadata.

Source	Description
JDBC, LDAP, or other types of datastore (if configured)	Select to evaluate attributes returned from a data source.
Mapped Attributes	Select to evaluate the mapped attributes.
Mapped from	Select to evaluate attributes from the authentication source.
Policy Contract IdP where	Visible and applicable only when configuring an access token mapping where the source of the attribute is something other than Client Credentials and Default . See .
Persistent Grant	Select to evaluate the default attribute USER_KEY and other extended attributes (if defined) from the persistent grant.
	Visible and applicable only when configuring an access token mapping where the source of the attribute is not Client Credentials .

- 3. From the **Attribute Name** list, select the attribute to be evaluated.
- 4. From the **Condition** list, select the comparison method.

Available methods:

- equal to
- equal to (case insensitive)
- equal to DN
- not equal to
- not equal to (case insensitive)
- not equal to DN
- multi-value contains
- multi-value contains (case insensitive)
- multi-value contains DN
- multi-value does not contain
- multi-value does not contain (case insensitive)
- multi-value does not contain DN



Note:

The first six conditions are intended for single-value attributes. Use one of the multi-value ... conditions for PingFederate to validate whether one of the attribute values matches the specified value. When an attribute has multiple values, using a single-value condition causes the criteria to fail.

5. In the Value field, enter the comparison value.



Values are compared verbatim. If you require complex evaluations, including conditional criteria or partial matching, define them using attribute mapping expressions. For more information, see Attribute mapping expressions on page 214.

6. In the **Error Result** field, enter a custom error message.

The value of this field is used by the error_description protocol field. Using an error code in the **Error Result** field allows an application to process the code in a variety of ways; for example, display an error message or e-mail an administrator.

To use localized descriptions, enter a unique alias in the Error Result field, such as someIssuanceCriterionFailed. Insert the same alias with the desired localized text in the applicable language resource files, located in the $<pf_install>/pingfederate/server/default/conf/language-packs directory.$

If not defined, PingFederate returns ACCESS DENIED when the criterion fails at runtime.

- 7. Click Add.
- 8. Optional: Repeat to add more criteria.
- If you require complex evaluations, including conditional criteria or partial matching, define them using attribute mapping expressions. For more information, see *Attribute mapping expressions* on page 214.
 - Click Show Advanced Criteria.
 - b. In the **Expression** field, enter the required expressions.
 - c. Optional: In the Error Result field, enter an error code or message.



Note:

If the expressions resolve to a string value instead of true or false, the returned value overrides the **Error Result** field value.

- d. Click Add.
- e. Optional: Click **Test**, enter values in the applicable fields, and verify the results.
- f. Optional: Repeat to add multiple criteria using attribute mapping expressions.

Reviewing the access token mapping

On the **Summary** tab, review your access token mapping configuration.

Steps

- To amend your configuration, click the corresponding tab title and then follow the steps to complete
 the task.
- To keep your changes, click Save.
- To discard your changes, click **Cancel**.

Configuring an OAuth assertion grant IdP connection

An OAuth assertion grant connection exchanges a SAML assertion or a JSON web token (JWT) for an OAuth access token with the PingFederate OAuth authorization server.

About this task

You can configure an OAuth assertion grant connection with an identity provider (IdP) partner either in conjunction with browser-based single sign-on (SSO), WS-Trust, or independently.

For more information, see Security Assertion Markup Language (SAML) 2.0 Profile for OAuth 2.0 Client Authentication and Authorization Grants and JSON Web Token (JWT) Profile for OAuth 2.0 Client Authentication and Authorization Grants.

Steps

1. Go to Authentication # Integration # IdP Connections and then click Create Connection.



You can also select other options, such as the Browser SSO Profiles check box. If you do, you will be prompted to complete the required configuration. This topic only focuses on the OAuth Assertion **Grant** configuration.

- 3. On the **General Info** tab, enter the required information.
- 4. On the OAuth Assertion Grant Attribute Mapping tab, click Configure OAuth Assertion Grant Attribute Mapping.

Defining an attribute contract for the OAuth assertion grant

About this task

An attribute contract is a set of user attributes the IdP sends in the SAML assertions or JWTs for this connection. You identity these attributes on the OAuth Assertion Grant Attribute Mapping # Attribute Contract window.

TOKEN SUBJECT represents the name identifier of the user for whom the access token is being requested, the SAML SUBJECT attribute in SAML assertions and the sub claim in JWTs.

Optionally, you can mask the values of attributes (other than **TOKEN SUBJECT**) in the log files that PingFederate writes when it receives security tokens.

Steps

- To add an attribute, follow these steps:
 - a. Enter the attribute name in the text box. Attribute names are case-sensitive and must correspond to the attribute names expected by your partner.
 - b. Select the check box under Mask Values in Log.
 - c. Click Add.
- To modify an attribute name or masking selection, follow these steps:
 - Click Edit under Action for the attribute.
 - b. Make the change and click **Update**.



Note:

If you change your mind, ensure that you click **Cancel** under **Action**.

To delete an attribute, click **Delete** under **Action** for the attribute.

Configuring access token manager mappings

To define how access tokens are created, use the Access Token Manager Mapping tab to associate one or more access token manager instances with this connection.

Steps

- 1. Within your IdP connection configuration, go to the Access Token Manager Mapping tab within the OAuth Assertion Grant Attribute Mapping tab.
- 2. Perform one of the following actions.

Action	Steps
Create a new Access Token Manager mapping configuration	Click Create New Access Token Manager Mapping.

Action	Steps
Edit an existing Access Token Manager mapping configuration	Click the mapping configuration's name.
Delete an Access Token Manager mapping	Click Delete under Action for the applicable mapping configuration.

Selecting an access token manager instance

This configuration maps attribute values from the (identity provider) IdP connection into the access token to define the resulting content of the access token.

Steps

- 1. Within your IdP connection configuration, go to the **Access Token Manager** tab.
 - a. Click the OAuth Assertion Grant Attribute Mapping tab, and then click Configure OAuth Assertion Grant Attribute Mapping.
 - b. Click the Access Token Manager Mapping tab, and then click Create New Access Token Manager Mapping.
- 2. From the Access Token Manager list, select an access token manager instance.



If the access token manager instance you need is not available, click Manage Access Token **Management Instances** to define one or more instances for this connection.

Configuring a datastore for OAuth assertion grant attribute mapping

You can optionally set up datastore queries to supplement values returned from the OAuth assertion grant attribute mapping source.

Steps

- 1. Within your identity provider (IdP) connection configuration, go to the **Data Store** tab.
 - a. Click the OAuth Assertion Grant Attribute Mapping tab, and then click Configure OAuth Assertion Grant Attribute Mapping.
 - b. Click the Access Token Manager Mapping tab, and then click Create New Access Token Manager Mapping.
 - c. Click the **Data Store** tab.
- 2. Perform one of the following actions.

Choose from:

- To set up datastore queries, select a datastore from the Active Data Store list, and then click Next. Follow the wizard to complete the setup. For configuration steps, see Datastore query configuration on page 362.
- To skip this optional configuration, select No Data Store and then click Next.

Configuring OAuth assertion grant contract fulfillment

Map values from the SAML assertions or JSON web tokens (JWTs) to the attributes defined for the attribute contract. The access token manager instance requires these values to create an OAuth access token.

About this task

At runtime, a single sign-on (SSO) operation fails if PingFederate cannot fulfill the required attribute.

Steps

 On OAuth Assertion Grant Attribute Mapping # OAuth Assertion Grant Attribute Mapping Configuration # Contract Fulfillment, select a source from the Source list and then choose or enter a value for each attribute.

Assertion

When selected, the Value list populates with attributes from the SAML assertion or the JWT.

For example, to map the value of **SAML_SUBJECT** from a SAML assertion, or **sub** from a JWT, as the value of an attribute on the access-token contract, select **Assertion** from the **Source** list and **TOKEN SUBJECT** from the **Value** list.

Context

When selected, the Value list populates with the available context of the transaction.



Note:

Because the **HTTP Request** context value is retrieved as a Java object rather than text, use OGNL expressions to evaluate and return values. For more information, see **Expression**.

Extended Client Metadata

Values are returned from the client record.

LDAP, JDBC, or Other

When selected, the **Value** list is populated with attributes that you have selected from the datastore. Select the desired attribute from the list.

Expression

When enabled, this option provides more complex mapping capabilities, such as transforming incoming values into different formats. Select **Expression** from the **Source** list, click **Edit** under **Actions**, and compose your OGNL expressions. All variables available for text entries are also available for expressions. For more information, see **Text**.

Expressions are not enabled by default. For more information about enabling and editing OGNL expressions, see .

No Mapping

Select this option to ignore the Value field.

Text

When selected, the text you enter is used at runtime. You can mix text with references to any of the values from the SSO token, using the \${attribute} syntax.

When applicable, you can enter values from your datastore using the \${ds.<attribute>} syntax, where <attribute> is any attribute that you have selected from the datastore.

2. Click Next.

Defining issuance criteria for OAuth assertion grants

Individual attributes within policy contracts can further determine whether PingFederate approves or rejects requests. You can define those criteria to satisfy or you can choose to skip this configuration.

About this task

On the **Issuance Criteria** tab, define the criteria to satisfy in order for PingFederate to further process a request. Use this token authorization feature to conditionally approve or reject requests based on individual attributes.

Begin this optional configuration by choosing the source that contains the attribute to verify. Some sources are common to almost all use cases, such as **Mapped Attributes**. Other sources depend on the type

You can define multiple criteria, which must all be satisfied in order for PingFederate to move a request to the next phase. A criterion is satisfied when the runtime value of the selected attribute matches or does not match the specified value, depending on the chosen comparison method. The multi-value contains ... or multi-value does not contain ... comparison methods are intended for attributes that can contain multiple values. Such a criterion is considered satisfied if one of the multiple values match or does not match the specified value. Values are compared verbatim. If you require complex evaluations, including conditional criteria or partial matching, define them using attribute mapping expressions.



Note:

All criteria defined must be satisfied, or evaluated as true, for a request to move forward, regardless of how the criteria were defined. As soon as one criterion fails, PingFederate rejects the request and returns an error message.

Steps

- 1. Within your identity provider (IdP) connection configuration, go to OAuth Assertion Grant Attribute Mapping # OAuth Assertion Grant Attribute Mapping Configuration # Issuance Criteria.
- 2. From the **Source** list, select the attribute's source.

Depending on the selection, the **Attribute Name** list populates with associated attributes. See the following table for more information.

Source	Description
Assertion	Select to evaluate attributes from the IdP connection.
Context	Select to evaluate properties returned from the context of the transaction at runtime.
	Note:
	As the HTTP Request context value is retrieved as a Java object rather than text, attribute mapping expressions are more appropriate to evaluate and return values.
Extended Client Metadata	Select to evaluate OAuth client metadata.
JDBC, LDAP, or other types of datastore (if configured)	Select to evaluate attributes returned from a data source.
Mapped Attributes	Select to evaluate the mapped attributes.

3. From the Attribute Name list, select the attribute to be evaluated.

4. From the **Condition** list, select the comparison method.

Available methods:

- equal to
- equal to (case insensitive)
- equal to DN
- not equal to
- not equal to (case insensitive)
- not equal to DN
- multi-value contains
- multi-value contains (case insensitive)
- multi-value contains DN
- multi-value does not contain
- multi-value does not contain (case insensitive)
- multi-value does not contain DN



Note:

The first six conditions are intended for single-value attributes. Use one of the multi-value ... conditions for PingFederate to validate whether one of the attribute values matches the specified value. When an attribute has multiple values, using a single-value condition causes the criteria to fail.

5. In the **Value** field, enter the comparison value.



Note:

Values are compared verbatim. If you require complex evaluations, including conditional criteria or partial matching, define them using attribute mapping expressions. For more information, see Attribute mapping expressions on page 214.

6. In the **Error Result** field, enter a custom error message.

The value of this field is used by the error description protocol field. Using an error code in the Error Result field allows an application to process the code in a variety of ways, such as displaying an error message or emailing an administrator.

To use localized descriptions, enter a unique alias in the Error Result field, such as someIssuanceCriterionFailed. Insert the same alias with the desired localized text in the applicable language resource files, located in the <pf install>/pingfederate/server/ default/conf/language-packs directory.

If not defined, PingFederate returns ACCESS DENIED when the criterion fails at runtime.

- 7. Click Add.
- 8. Optional: Repeat to add more criteria.

- 9. If you require complex evaluations, including conditional criteria or partial matching, define them using attribute mapping expressions. For more information, see Attribute mapping expressions on page 214.
 - a. Click Show Advanced Criteria.
 - b. In the **Expression** field, enter the required expressions.
 - c. Optional: In the **Error Result** field, enter an error code or message.



Note:

If the expressions resolve to a string value instead of true or false, the returned value overrides the Error Result field value.

- d. Click Add.
- e. Optional: Click **Test**, enter values in the applicable fields, and verify the results.
- f. Optional: Repeat to add multiple criteria using attribute mapping expressions.

Reviewing OAuth assertion grant attribute mapping configuration

On the **Summary** tab, review your OAuth assertion grant attribute mapping configuration.

Steps

- To amend your configuration, click the corresponding tab title, then follow the configuration wizard to complete the task.
- To keep your changes, click **Done** and continue with the rest of the configuration.



Tip:

When editing an existing configuration, you can also click **Save** as soon as the administrative console offers the opportunity to do so.

To discard your changes, click Cancel.

Reviewing OAuth assertion grant configuration

On the **Summary** tab, review your OAuth assertion grant configuration.

Steps

- To amend your configuration, click the corresponding tab title, then follow the configuration wizard to complete the task.
- To keep your changes, click **Done** and continue with the rest of the configuration.



When editing an existing configuration, you can also click **Save** as soon as the administrative console offers the opportunity to do so.

To discard your changes, click Cancel.

Configuring OpenID Connect policies

This configuration allows you to define OpenID Connect policies for client access to attributes mapped according to OpenID specifications.

About this task

You can include a session identifier in the ID Tokens, which might be useful for the relying parties, such as PingAccess, for client session management.

Go to Applications # OAuth # OpenId Connect Policy Management and perform one of the following actions.

Action	Steps
Configure a new OpenID Connect policy	Click Add Policy.
Modify an existing OpenID Connect policy	Select its name under Policy ID .
Review the usage of an existing OpenID Connect policy	Click Check Usage under Action.
Remove an existing OpenID Connect policy or to cancel the removal request	Click Delete or Undelete under Action .
Elect an existing OpenID Connect policy to be the default OpenID Connect policy	Click Set as Default under Action.

Configuring policy and ID token settings

On the Manage Policy tab, enter the required information and configure optional settings for ID tokens issued under this policy.

Steps

- 1. Go to Applications # OAuth # OpenID Connect Policy Management and click Add Policy.
- 2. In the Policy ID field, enter the policy identifier.
- 3. In the **Name** field, enter the policy name.
- 4. From the Access Token Manager list, select an access token management instance.
- 5. Optional: In minutes, define the expiry information for ID tokens issued based on this policy in the ID Token Lifetime field.

The default value is 5 minutes.

- 6. Optional: Select the Include Session Identifier in ID Token check box to add a session identifier (pi.sri) in the ID tokens.
- 7. Optional: Select the Include User Info in ID Token check box to include additional attributes in the ID tokens.



OAuth clients can also obtain additional attributes from the UserInfo endpoint at /idp/ userinfo.openid. For more information, see *UserInfo endpoint* on page 1085.

8. Optional: Select the Include State Hash in ID Token check box to include the s hash claim in ID tokens.



Note:

A state hash protects the state parameter by binding it to the ID token. For more information, see Financial Services - Financial API - Part 2: Read and Write API Security Profile.

9. Optional: Select the Return ID Token On Refresh Grant checkbox to return an ID token for OpenID Connect to Salesforce and Kubernetes when the OAuth access token is refreshed.

In the Attribute Contract tab, you can define the list of attributes that PingFederate can return to the OAuth clients.

About this task

Every new OpenID Connect policy contract begins with a list of standard attributes. These attributes or claims are defined in the OpenID Connect specification. You can optionally remove standard attributes, turn them into non-standard attributes, or add new non-standard attributes.



Note:

In OpenID Connect, scopes affect the list of attributes that PingFederate can return to the OAuth clients. The attributes that PingFederate returns to OAuth clients vary, depending on the scopes originally approved by the resource owner.

By default, all attributes defined on this window are deliverable through the UserInfo endpoint. If an implicit client makes a token request by providing id token as the only response type parameter value, the client will only receive an ID token without an access token. As the client will not be able to retrieve additional attributes from the UserInfo endpoint without a valid access token, PingFederate includes the applicable attributes in the ID token instead.

If you have not selected the Include User Info in ID Token option in the Manage Policy tab for this policy, you can choose how attributes are delivered to clients. Similar to the default delivery behavior, in the scenario where an implicit client makes a token request by providing id token as the sole response type parameter value, PingFederate includes the applicable attributes in the ID token regardless of any configured overrides.

Steps

- To add a new attribute:
 - a. Enter the name of the attribute under **Extend the Contract**.
 - b. Optional: Select the Override Default Delivery check box to choose how the attribute is
 - Select the check box under ID Token if this attribute can be included in ID tokens.
 - Select the check box under UserInfo if this attribute can be included in UserInfo responses.
 - c. Click Add.
- To modify an existing entry, use the Edit, Update, and Cancel buttons. Choose how the attribute is delivered, as needed.
- To remove an existing entry, click **Delete**.

Configuring attribute scopes

With OpenID Connect, scopes affect the list of attributes that PingFederate can return to the OAuth clients. In the Attribute Scopes tab, you can optionally add associations between scopes and attributes beyond what is defined in the specification.

Steps

1. Go to Applications # OAuth # OpenID Connect Policy Management and select your policy, or click Add Policy.

- 2. In the Attribute Scopes tab, add any number of scope-to-attributes associations.
 - a. Select a scope from the **Scope** list.

Common and exclusive scopes are both available.

b. Select the relevant check boxes under Attributes.



Note:

If you have selected a standard scope, its associated standard attributes are automatically selected and cannot be modified. You can select additional attributes to be associated with the selected scope.

If you have selected the profile scope, any non-standard attributes that are not associated with the profile scope become inaccessible to your OAuth clients. The administrative console displays a warning message with a list of inaccessible attributes. Select the relevant check boxes to make the non-standard attributes accessible, or ignore the message to leave them inaccessible for now.

- c. Click Add.
- Optional: Repeat these steps to define additional scope-to-attributes associations.

Click **Edit**, **Update**, or **Cancel** to make or undo a change to an existing entry. Click **Delete** or **Undelete** to remove an existing entry or cancel the removal request.

3. Click Next.

Configuring policy attribute sources and user lookup

You can optionally set up datastore queries to supplement values returned from the policy attribute source.

Steps

- Go to Applications # OAuth # OpenID Connect Policy Management and select your policy, or click Add New Policy.
- On the Attribute Sources & User Lookup tab, perform one of the following actions. Choose from:
 - To set up datastore queries, click Add Attribute Source. Complete the setup on the Attribute Sources & User Lookup tab. For configuration steps, see Datastore query configuration on page 362.
 - To skip this optional configuration, click **Next**.

Configuring ID token fulfillment

On the **Contract Fulfillment** tab, map attributes from the access token or other sources to fulfill the attribute contract.

Steps

1. Go to Applications # OAuth # OpenID Connect Policy Management and select your policy, or click Add Policy.

2. On the Contract Fulfillment tab, select a source from the Source list and then select or enter a value for each attribute in the contract.

Map the subject attribute and all extended attributes from one of the following sources:

Context

Values are returned from the context of the transaction at runtime.



As the HTTP Request context value is retrieved as a Java object rather than text, OGNL expressions are preferred to evaluate and return values.

To enter an expression, select **Expression** under **Source**, and then click **Edit**.

If **Expression** is not available, you can enable it by editing the

org.sourceid.common.ExpressionManager.xml file in the <pf install>/ pingfederate/server/default/data/config-store directory.

Extended Client Metadata

Values are returned from the client record.

LDAP/JDBC/Other

Values are returned from your datastore, if used. When selected, the **Value** list populates with attributes from the datastore.

Expression

When enabled, this option provides more complex mapping capabilities, such as transforming incoming values into different formats. All of the variables available for text entries are also available for expressions.

No Mapping

This option ignores the Value field.

Text

The value is what you enter. This can be text only, or you can mix text with references to the unique user ID returned from the credentials validator, using the \${attribute} syntax.

You can also enter values from your datastore, when applicable, using the \${ds.attribute} syntax, where attribute is any of the datastore attributes you have selected.

Access Token

The value is provided from the access token.

Persistent Grant

Enables direct mapping from the grant to the ID Token and to user information attributes.

3. Click Next.

Defining issuance criteria for policy mapping

Individual attributes within policy contracts can further determine whether PingFederate approves or rejects requests. You can define those criteria to satisfy or you can choose to skip this configuration.

About this task

On the Issuance Criteria tab, define the criteria to satisfy in order for PingFederate to further process a request. Use this token authorization feature to conditionally approve or reject requests based on individual attributes.

Begin this optional configuration by choosing the source that contains the attribute to verify. Some sources are common to almost all use cases, such as Mapped Attributes. Other sources depend on the type of configuration, such as JDBC. Irrelevant sources are automatically hidden. Once you select a source,

choose the attribute to verify. Depending on the selected source, the available attributes or properties vary. Specify the comparison condition and the desired value to compare to.

You can define multiple criteria, which must all be satisfied in order for PingFederate to move a request to the next phase. A criterion is satisfied when the runtime value of the selected attribute matches or does not match the specified value, depending on the chosen comparison method. The multi-value contains ... or multi-value does not contain ... comparison methods are intended for attributes that can contain multiple values. Such a criterion is considered satisfied if one of the multiple values match or does not match the specified value. Values are compared verbatim. If you require complex evaluations, including conditional criteria or partial matching, define them using attribute mapping expressions.



All criteria defined must be satisfied, or evaluated as true, for a request to move forward, regardless of how the criteria were defined. As soon as one criterion fails, PingFederate rejects the request and returns an error message.

Steps

- 1. Go to Applications # OAuth # OpenID Connect Policy Management and select your policy, or click Add Policy.
- 2. On the Issuance Criteria tab, select the attribute's source from the Source list.

Depending on the selection, the Attribute Name list populates with associated attributes. See the following table for more information.

Source	Description	
Access Token	Select to evaluate attributes from the access token.	
Context	Select to evaluate properties returned from the context of the transaction at runtime.	
	Note: The HTTP Request context value is retrieved as a Java object rather than text. For this reason, attribute mapping expressions are more appropriate to evaluate and return values.	
Extended Client Metadata	Select to evaluate OAuth client metadata.	
JDBC, LDAP, or other types of datastore (if configured)	Select to evaluate attributes returned from a data source.	
Mapped Attributes	Select to evaluate the mapped attributes.	
Persistent Grant	Select to evaluate attributes from the persistent grant.	

3. From the **Attribute Name** list, select the attribute to be evaluated.

Available methods:

- equal to
- equal to (case insensitive)
- equal to DN
- not equal to
- not equal to (case insensitive)
- not equal to DN
- multi-value contains
- multi-value contains (case insensitive)
- multi-value contains DN
- multi-value does not contain
- multi-value does not contain (case insensitive)
- multi-value does not contain DN



Note:

The first six conditions are intended for single-value attributes. Use one of the multi-value ... conditions for PingFederate to validate whether one of the attribute values matches the specified value. When an attribute has multiple values, using a single-value condition causes the criteria to fail.

5. In the **Value** field, enter the comparison value.



Note:

Values are compared verbatim. If you require complex evaluations, including conditional criteria or partial matching, define them using attribute mapping expressions. For more information, see Attribute mapping expressions on page 214.

6. In the **Error Result** field, enter a custom error message.

The value of this field is used by the error description protocol field. Using an error code in the Error Result field allows an application to process the code in a variety of ways, such as displaying an error message or e-mailing an administrator.

To use localized descriptions, enter a unique alias in the Error Result field, such as someIssuanceCriterionFailed. Insert the same alias with the desired localized text in the applicable language resource files, located in the <pf install>/pingfederate/server/ default/conf/language-packs directory.

If not defined, PingFederate returns ACCESS DENIED when the criterion fails at runtime.

- 7. Click Add.
- 8. Optional: Repeat to add more criteria.

- 9. If you require complex evaluations, including conditional criteria or partial matching, define them using attribute mapping expressions. For more information, see Attribute mapping expressions on page 214.
 - a. Click Show Advanced Criteria.
 - b. In the **Expression** field, enter the required expressions.
 - c. Optional: In the **Error Result** field, enter an error code or message.



Note:

If the expressions resolve to a string value instead of true or false, the returned value overrides the Error Result field value.

- d. Click Add.
- e. Optional: Click **Test**, enter values in the applicable fields, and verify the results.
- f. Optional: Repeat to add multiple criteria using attribute mapping expressions.

Reviewing your OpenID Connect policy

On the **Summary** tab, review your OpenID Connect policy.

Steps

- To amend your configuration, click the corresponding tab title and then follow the steps to complete the task.
- To keep your changes, click Save.
- To discard your changes, click Cancel.

Result



Note:

If this is the first policy you are creating, you must click Done and designate this policy as the default before saving on the OpenID Connect Policy Management window. You can change the default when you create additional policies.

Client Initiated Backchannel Authentication (CIBA)

Client Initiated Backchannel Authentication is an extension to OpenID Connect that improves the end-user experience during authentication and authorization in a federated environment.

The CIBA extension defines a new OAuth grant type where user consent can be requested through an out-of-band flow. CIBA improves the user experience, such as when making an online purchase from a merchant, because it does not require a browser redirect to a financial institution to authorize the purchase. Instead, the user can receive a push notification sent to the financial institution's native mobile app running on the user's phone to complete the authorization. For more information, see Configuring a CIBA authenticator instance.



Note:

The MFA Integration Kit includes the MFA CIBA Authenticator, which works with 's CIBA feature. For instructions on configuring the MFA CIBA Authenticator, see Configuring a CIBA authenticator instance.

A CIBA configuration consists of two components: a CIBA authenticator and a CIBA request policy.

CIBA authenticator

A CIBA authenticator is responsible for authenticating users through an out-of-band method.

You can use the PingFederate SDK to implement a custom solution. For more information, see the Javadoc for the OOBAuthPlugin interface, the SampleEmailAuthPlugin, java file for a sample implementation, and the SDK developer's guide for build and deployment information.

Once deployed, you can create one or more instance configurations of the authenticator.

For more information, see Configuring a CIBA authenticator instance on page 581.

CIBA request policy

CIBA request policies process identity hints and authenticate users to receive consent. Each request policy is associated with an instance of a CIBA authenticator. The CIBA grant flow is initiated by a direct request from the client and involves an out-of-band interaction with the user to complete authentication and authorization. OAuth clients that support the CIBA grant type can be configured to use a specific CIBA request policy or a default.

For more information, see *Defining a request policy* on page 582.



Because the CIBA extension is an OAuth grant type, to enable CIBA for the client, you must select CIBA in the **Allowed Grant Types** setting. Once selected, you can configure more client CIBA-related settings.

For more information, see Configuring an OAuth client on page 515.

Managing CIBA authenticators

Manage the Client Initiated Backchannel Authentication (CIBA) authenticators in PingFederate.

About this task

A CIBA authenticator is responsible for authenticating users through an out-of-band method.

You can use the PingFederate SDK to implement a custom solution. For more information, see the Javadoc for the OOBAuthPlugin interface, the SampleEmailAuthPlugin.java file for a sample implementation, and the SDK developer's guide for build and deployment information.



The Javadoc for PingFederate and the sample implementation are in the <pf install>/ pingfederate/sdk directory.

Once deployed, you can create one or more instance configurations of the authenticator.

Steps

- Go to Authentication # OAuth # CIBA Authenticators and create or select the instance you want to manage.
- To manage the CIBA authenticator, choose from the following options.

Option	Description	
Create New Instance	Configure a new instance	
<existing instance="" link=""> under Instance Name</existing>	me Modify an existing instance	
Check Usage	Review the usage of an existing instance	
Delete or Undelete	Remove an existing instance or cancel the removal	

Configuring a CIBA authenticator instance

About this task

The MFA Integration Kit includes the SDK CIBA Authenticator, which works with 's CIBA feature.



Note:

For instructions on configuring the CIBA Authenticator for the MFA Integration Kit, see Configuring a CIBA authenticator instance.

Steps

- 1. Go to Authentication # OAuth # CIBA Authenticators to open the CIBA Authenticators window.
- 2. On the CIBA Authenticators window, click Create New Instance to start the Create CIBA Authenticator Instance configuration workflow.
- 3. On the **Type** tab, configure the basics of this authenticator instance.
 - a. Enter a name and an ID in the Instance Name and Instance ID fields.
 - b. Select a CIBA authenticator from the **Type** list.

Selections vary depending on the deployed CIBA authenticators.

You can use the PingFederate SDK to implement a custom solution. For more information, see the Javadoc for the OOBAuthPlugin interface, the SampleEmailAuthPlugin.java file for a sample implementation and the SDK developer's guide for build and deployment information.



The Javadoc for PingFederate and the sample implementation are in the <pf install>/ pingfederate/sdk directory.

- 4. On the Instance Configuration tab, follow the on-screen instructions to configure the authenticator instance.
 - Configuration requirements vary depending on the authenticator implementation.
- 5. On the **Actions** tab, follow the on-screen instructions to test the validity of the authenticator instance configuration or to perform secondary configuration tasks.
 - Availability of this tab and actions vary depending on the authenticator implementation.
- 6. On the **Extended Contracts** tab, follow the on-screen instructions to define additional attributes.

The authenticator contract is the list of input parameters used to challenge the user for authentication. Some authenticators support extending the contract for additional functionality, such as formatting the data presented to the user during the authentication challenge.

- Availability of this window and supported attribute names vary depending on the authenticator implementation.
- 7. On the **Summary** tab, review your configuration, modify as needed, and click **Done** to exit the **Create** CIBA Authenticator Instance workflow.
- 8. On the CIBA Authenticators window, click Save to retain the configuration of the authenticator instance.

If you want to exit without saving the configuration, click Cancel.

Managing CIBA request policies

You can configure, modify, review, remove, and elect existing client initiated backchannel authentication (CIBA) request policies in the administrative console.

About this task

CIBA request policies process identity hints and authenticate users to receive consent. Each request policy is associated with an instance of a CIBA authenticator. The CIBA grant flow is initiated by a direct request from the client and involves an out-of-band interaction with the user to complete authentication and authorization. OAuth clients that support the CIBA grant type can be configured to use a specific CIBA request policy or a default.

Steps

- Go to Applications # OAuth # CIBA Request Policies. To configure a new CIBA request policy, click Add Policy.
- 2. Select an action from the following options:

Choose from:

- To modify an existing CIBA request policy, select it by its name under **Policy ID**.
- To review the usage of an existing CIBA request policy, click **Check Usage** under **Action**.
- To remove an existing CIBA request policy or to cancel the removal request, click **Delete** or Undelete under Action.
- To elect an existing CIBA request policy to be the default CIBA request policy, click Set as Default under Action.

Defining a request policy

You can define the basics of your client-initiated backchannel authentication (CIBA) request policy in the PingFederate administrative console.

Steps

- 1. Go to Applications # OAuth # CIBA Request Policies.
- 2. On the **Manage Policy** tab, define the basics of your CIBA request policy.

For more information about each field, refer to the following table.

Field	Description	
Policy ID	The unique identifier of this request policy.	
(Required)		
Name	The name of this request policy.	
(Required)		
Authenticator	The CIBA authenticator instance associated with this request policy.	
(Required)		

Field	Description	
User Code PCV	The Password Credential Validator (PCV) instance that PingFederate uses to validate the user_code parameter values it receives from clients associated with this request policy.	
	Important:	
	If a client is associated with a request policy that has been configured with a PCV instance, it can support user code in its configuration.	
	A client supporting user code must not be associated with a request policy that is not configured with a PCV instance. For more information on CIBA client configuration, see <i>Configuring an OAuth client</i> on page 515.	
Transaction Lifetime (Seconds)	The validity, in seconds, of authentication requests PingFederate receives from clients associated with this request policy since the generation of their authentication request acknowledgments.	
	The default value is 120.	
	Clients can request a shorter lifetime by including the requested_expiry request parameter in their authentication requests.	
Allow Unsigned Login Hint Token	Controls whether clients associated with this request policy can use unsigned JSON web tokens (JWT) as values of the <code>login_hint_token</code> request parameter in their authentication requests.	
	This check box is not selected by default.	
Require Token for Identity Hint	Controls whether clients associated with this request policy must use either the id_token_hint or login_hint_token as the identity hint in their authentication requests.	
	This check box is not selected by default.	
	When selected, clients associated with this request policy cannot use login_hint as the identity hint in their authentication requests.	
Alternative Login Hint Token Issuers	Alternative issuers that clients associated with this request policy can use in their signed login hint tokens. Furthermore, each additional issue requires either the JWKS url or the actual JWKS so that PingFederate c verify the authenticity of the signed login hint tokens.	

3. Click Next.

Configuring identity hint contract

You can configure the identity hint contract, which contains the set of attributes received in the client initiated backchannel authentication (CIBA) request that identifies the user.

About this task

IDENTITY HINT SUBJECT is a core attribute and is automatically populated by the sub attribute of an identity hint token, if found, or the attribute value of the login hint request attribute.

A client can send an ID token, id_token_hint, or a login hint token, login_hint_token, as the identity hint token. If you extend the identity hint contract with attribute names from the identity token, PingFederate fulfills them with values found in the identity token.



As needed, all attributes can optionally be fulfilled differently on the Identity Hint Contract Fulfillment tab.

Steps

- 1. Optional: Go to Applications # OAuth # CIBA Request Policies. On the Identity Hint Contract Fulfillment tab, enter an attribute name under Extend the Contract, and then click Add.
- 2. Repeat the previous step to define additional attributes. Click Next.

Use the **Edit**, **Update**, and **Cancel** workflow to make or undo a change to an existing entry. Click **Delete** to remove an entry.

Example

Example

Suppose the following JSON web token (JWT) matches the expected structure of the login hint tokens.

```
{
   "sub": "asmith",
   "attrs": {
      "mail": "asmith@example.com",
      "phone": "555-555-5555"
   }
}
```

To add both the mail and phone attributes, extend the contract with login_hint_token.attrs.mail and login hint token.attrs.phone, respectively.

Configuring identity hint contract fulfillment

You can process the identity hint to further augment the identity data prior to contract fulfillment in the PingFederate administrative console.

Steps

- 1. Go to Applications # OAuth # CIBA Request Policies # Identity Hint Contract Fulfillment.
- 2. Click **Manage Fulfillment** to begin the mapping configuration.
- 3. When the administrative console returns you to the Identity Hint Contract Fulfillment tab, click Next.

Configuring attribute sources and user lookup

You can configure the attribute sources and user lookup values in the PingFederate administrative console.

About this task

You can set up datastore queries to supplement values returned from the source. This configuration to fulfill the identity hint's attribute contract is optional.

Steps

1. Go to Applications # OAuth # CIBA Request Policies.

2. Choose one of the following actions:

Choose from:

- On the Attribute Sources & User Lookup tab, to set up datastore queries, click Add Attribute Source.
- To skip this optional configuration, click **Next**.

Follow the configuration workflow to complete the setup. For configuration steps, see *Datastore query configuration* on page 362.

When the administrative console returns you to the **Attribute Sources & User Lookup** tab, click **Next**.

Fulfilling identity hint contract

You can fulfill identity hint contracts in PingFederate.

About this task

On the Identity Hint Contract Fulfillment tab, fulfill the identity hint contract with values from the original identity hint, datastores, dynamic text values, or attribute mapping expressions, if enabled.

Steps

1. From the **Source** list, select a source.

For more information about the **Source** list, see the following table.

Source	Description	
Context	Select Context to return specific information from the request.	
JDBC, LDAP, or other types of datastore (if	Select an attribute source when PingFederate should retrieve attribute value from a datastore.	
configured)	When you make this selection, the list under Value populates with attributes from your database, directory, or other datastore.	
	Applicable only if you have added at least one attribute source on the Attribute Sources & User Lookup tab. For more information, see <i>Configuring attribute sources and user lookup</i> on page 584.	
Request	Select Request to use the attribute value PingFederate found in the CIBA request without customization.	
Expression (if enabled)	Select Expression to support complex mapping requirements, such as transforming incoming values into different formats. Additionally, HTTP request is retrieved as a Java object rather than text. For this reason, select Expression as the source and use OGNL expressions to evaluate and return specific information from the HTTP request.	
	Applicable only if you have enabled the use of expressions in PingFederate. For more information, see <i>Attribute mapping expressions</i> on page 214.	
No Mapping	Select No Mapping to ignore the Value field, making value selection unnecessary.	

Source	Description
Text	Select Text to return the value you enter under Value .
	You might use a static text value if the target web application provides a service based on the name of your organization.
	You can mix text with references to attributes from the identity provider (IdP) adapter contract by using the \${ <attribute>} syntax.</attribute>
	\${ds. <attr-source-id>.<attribute>}You can also enter references to syntax, where <attr-source-id> is the Attribute Source ID value you entered on the Data Store tab and <attribute> is an attribute from the datastore.</attribute></attr-source-id></attribute></attr-source-id>

2. Specify a value associated with the selected source.

Not applicable if you selected Request. You can also enter references to attributes from configured attribute sources by using the Source list.

- 3. Repeat these steps until all attributes are configured.
- 4. Click Next.

Defining issuance criteria for identity hint contract

You must define issuance criteria for an identity hint contract to further process a request in PingFederate.

About this task

On the Issuance Criteria tab, define the criteria to satisfy in order for PingFederate to further process a request. Use this token authorization feature to conditionally approve or reject requests based on individual attributes.

Begin this optional configuration by choosing the source that contains the attribute to verify. Some sources are common to almost all use cases, such as Mapped Attributes. Other sources depend on the type of configuration, such as **JDBC**. Irrelevant sources are automatically hidden. Once you select a source, choose the attribute to verify. Depending on the selected source, the available attributes or properties vary. Specify the comparison condition and the desired value to compare to.

You can define multiple criteria, which must all be satisfied in order for PingFederate to move a request to the next phase. A criterion is satisfied when the runtime value of the selected attribute matches or does not match the specified value, depending on the chosen comparison method. The multi-value contains ... or multi-value does not contain ... comparison methods are intended for attributes that can contain multiple values. Such a criterion is considered satisfied if one of the multiple values match or does not match the specified value. Values are compared verbatim. If you require complex evaluations, including conditional criteria or partial matching, define them using attribute mapping expressions.



Note:

All criteria defined must be satisfied, or evaluated as true, for a request to move forward, regardless of how the criteria were defined. As soon as one criterion fails, PingFederate rejects the request and returns an error message.

1. From the **Source** list, select the attribute's source.

Depending on the selection, the Attribute Name list populates with associated attributes. See the following table for more information.

Source	Description	
Context	Select to evaluate properties returned from the context of the transaction at runtime.	
	Note:	
	As the HTTP Request context value is retrieved as a Java object rather than text, attribute mapping expressions are more appropriate to evaluate and return values.	
JDBC, LDAP, or other types of datastore (if configured)	Select to evaluate attributes returned from a data source.	
Mapped Attributes	Select to evaluate the mapped attributes.	
Request	Select to evaluate attributes from the CIBA request.	

- 2. From the **Attribute Name** list, select the attribute to be evaluated.
- 3. From the **Condition** list, select the comparison method.

Available methods:

- equal to
- equal to (case insensitive)
- equal to DN
- not equal to
- not equal to (case insensitive)
- not equal to DN
- multi-value contains
- multi-value contains (case insensitive)
- multi-value contains DN
- multi-value does not contain
- multi-value does not contain (case insensitive)
- multi-value does not contain DN



Note:

The first six conditions are intended for single-value attributes. Use one of the multi-value ... conditions for PingFederate to validate whether one of the attribute values matches the specified value. When an attribute has multiple values, using a single-value condition causes the criteria to fail.

4. In the Value field, enter the comparison value.



Note:

Values are compared verbatim. If you require complex evaluations, including conditional criteria or partial matching, define them using attribute mapping expressions. For more information, see Attribute mapping expressions on page 214.

5. In the **Error Result** field, enter a custom error message.

To use localized descriptions, enter a unique alias in the Error Result field, such as someIssuanceCriterionFailed. Insert the same alias with the desired localized text in the applicable language resource files, located in the $<pf_install>/pingfederate/server/default/conf/language-packs directory.$

If not defined, PingFederate returns ACCESS DENIED when the criterion fails at runtime.

- 6. Click Add.
- 7. Optional: Repeat to add more criteria.
- 8. If you require complex evaluations, including conditional criteria or partial matching, define them using attribute mapping expressions. For more information, see *Attribute mapping expressions* on page 214.
 - a. Click Show Advanced Criteria.
 - b. In the **Expression** field, enter the required expressions.
 - c. Optional: In the Error Result field, enter an error code or message.



Note

If the expressions resolve to a string value instead of true or false, the returned value overrides the **Error Result** field value.

- d. Click Add.
- e. Optional: Click Test, enter values in the applicable fields, and verify the results.
- f. Optional: Repeat to add multiple criteria using attribute mapping expressions.
- 9. Click Next.

Reviewing identity hint contract fulfillment

You can review, modify, or discard identity hint contract fulfillment configurations in PingFederate.

About this task

On the **Summary** tab, review your configuration.

Steps

Perform the following actions as needed.

Action	How to accomplish it	
	Click the corresponding tab and follow the configuration workflow	
Keep your changes	Click Done and continue with the configuration	
Discard your changes	Click Cancel	

Configuring attribute sources and user lookup for request policy contract

You can configure the attribute sources and user lookup for Client-Initiated Backchannel Authentication (CIBA) request policy contracts in the PingFederate administrative console.

About this task

You can optionally set up datastore queries to supplement values returned from the source. This configuration to fulfill the request policy's attribute contract is optional.

Steps

1. Go to Applications # OAuth # CIBA Request Policies.

Choose from:

- On the Attribute Sources & User Lookup tab, to set up datastore queries, click Add Attribute Source.
- To skip this optional configuration, click **Next**.

Follow the configuration workflow to complete the setup. For configuration steps, see *Datastore query configuration* on page 362.

When the administrative console returns you to the **Attribute Sources & User Lookup** tab, click **Next**.

Configuring request policy contract fulfillment

You can fulfill the request policy contract in PingFederate.

About this task

On the **Contract Fulfillment** tab, fulfill the request policy contract with values from the original identity hint, datastores, dynamic text values, or attribute mapping expressions (if enabled).

This contract is used to map into the OAuth grant (the USER_KEY attribute), the Client Initiated Back channel Authentication (CIBA) authenticator (attributes vary depending on the authenticator), and the user code Password Credential Validator (PCV) (the USER_CODE_USER_NAME attribute). The USER_CODE_USER_NAME attribute is shown only if a PCV instance is selected on the Manage Policy window.

Steps

1. Select a source from the **Source** list.

For more information about the **Source** list, see the following table.

Source	Description	
Context	Select Context to return specific information from the request.	
JDBC, LDAP, or other types of datastores (if	Select an attribute source when PingFederate should retrieve attribute value from a datastore.	
configured)	When you make this selection, the list under Value is populated with attributes from your database, directory, or other datastore.	
	Applicable only if you have added at least one attribute source on the Attribute Sources & User Lookup window. For more information, see <i>Configuring attribute sources and user lookup for request policy contract</i> on page 588.	
Request	Select Request to use the attribute value PingFederate found in the CIBA request without customization.	
Expression (if enabled)	Supports complex mapping requirements, such as transforming incoming values into different formats. Additionally, the HTTP request is retrieved as a Java object rather than text. Therefore, select Expression as the source and use OGNL expressions to evaluate and return specific information from the HTTP request.	
	Applicable only if you have enabled the use of expressions in PingFederate. For more information, see <i>Attribute mapping expressions</i> on page 214.	
No Mapping	Select No Mapping to ignore the Value field, making value selection unnecessary.	

Source	Description
Text	Select Text to return the value you entered under Value .
	You might use a static text value if the target web application provides a service based on the name of your organization. You can provide the attribute value as a constant.
	You can mix text with references to attributes from the identity provider (IdP) adapter contract by using the \${ <attribute>} syntax.</attribute>
sources by using the syntax, where <at: <b="" entered="" on="" the="">Attrik</at:>	You can also enter references to attributes from configured attribute sources by using the \${ds. <attr-source-id>.<attribute>} syntax, where <attr-source-id> is the Attribute Source ID value you entered on the Attribute Sources & User Lookup # Data Store tab and <attribute> is an attribute from datastore.</attribute></attr-source-id></attribute></attr-source-id>

- 2. Specify a value associated with the selected source.
- 3. Repeat these steps until all attributes are configured.
- 4. Click Next.

Defining issuance criteria for CIBA request policy

You must define the issuance criteria for PingFederate to further process a client-initiated backchannel authentication (CIBA) request policy.

About this task

On the Issuance Criteria tab, define the criteria to satisfy in order for PingFederate to further process a request. Use this token authorization feature to conditionally approve or reject requests based on individual attributes.

Begin this optional configuration by choosing the source that contains the attribute to verify. Some sources are common to almost all use cases, such as Mapped Attributes. Other sources depend on the type of configuration, such as **JDBC**. Irrelevant sources are automatically hidden. Once you select a source, choose the attribute to verify. Depending on the selected source, the available attributes or properties vary. Specify the comparison condition and the desired value to compare to.

You can define multiple criteria, which must all be satisfied in order for PingFederate to move a request to the next phase. A criterion is satisfied when the runtime value of the selected attribute matches or does not match the specified value, depending on the chosen comparison method. The multi-value contains ... or multi-value does not contain ... comparison methods are intended for attributes that can contain multiple values. Such a criterion is considered satisfied if one of the multiple values match or does not match the specified value. Values are compared verbatim. If you require complex evaluations, including conditional criteria or partial matching, define them using attribute mapping expressions.



All criteria defined must be satisfied, or evaluated as true, for a request to move forward, regardless of how the criteria were defined. As soon as one criterion fails, PingFederate rejects the request and returns an error message.

1. From the **Source** list, select the attribute's source.

Depending on the selection, the Attribute Name list populates with associated attributes. See the following table for more information.

Source	Description	
Context	Select to evaluate properties returned from the context of the transaction at runtime.	
	Note:	
	The HTTP Request context value is retrieved as a Java object rather than text. For this reason, attribute mapping expressions are more appropriate to evaluate and return values.	
JDBC, LDAP, or other types of datastore (if configured)		
Mapped Attributes	Select to evaluate the mapped attributes.	
Request	Select to evaluate attributes from the CIBA request.	

- 2. From the **Attribute Name** list, select the attribute to be evaluated.
- 3. From the **Condition** list, select the comparison method.

Available methods:

- equal to
- equal to (case insensitive)
- equal to DN
- not equal to
- not equal to (case insensitive)
- not equal to DN
- multi-value contains
- multi-value contains (case insensitive)
- multi-value contains DN
- multi-value does not contain
- multi-value does not contain (case insensitive)
- multi-value does not contain DN



Note:

The first six conditions are intended for single-value attributes. Use one of the multi-value ... conditions for PingFederate to validate whether one of the attribute values matches the specified value. When an attribute has multiple values, using a single-value condition causes the criteria to fail.

4. In the Value field, enter the comparison value.



Note:

Values are compared verbatim. If you require complex evaluations, including conditional criteria or partial matching, define them using attribute mapping expressions. For more information, see Attribute mapping expressions on page 214.

5. In the **Error Result** field, enter a custom error message.

To use localized descriptions, enter a unique alias in the Error Result field, such as someIssuanceCriterionFailed. Insert the same alias with the desired localized text in the applicable language resource files, located in the <pf install>/pingfederate/server/ default/conf/language-packs directory.

If not defined, PingFederate returns ACCESS DENIED when the criterion fails at runtime.

- 6. Click Add.
- 7. Optional: Repeat to add more criteria.
- 8. If you require complex evaluations, including conditional criteria or partial matching, define them using attribute mapping expressions. For more information, see Attribute mapping expressions on page 214.
 - a. Click Show Advanced Criteria.
 - b. In the **Expression** field, enter the required expressions.
 - c. Optional: In the Error Result field, enter an error code or message.



If the expressions resolve to a string value instead of true or false, the returned value overrides the Error Result field value.

- d. Click Add.
- e. Optional: Click **Test**, enter values in the applicable fields, and verify the results.
- f. Optional: Repeat to add multiple criteria using attribute mapping expressions.
- 9. Click Next.

Reviewing your CIBA request policy

Client Initiated Backchannel Authentication (CIBA) is an extension to OpenID Connect that can improve the end-user experience during authentication and authorization in a federated environment. You can review the request policy in PingFederate.

Steps

- To amend your configuration, click the corresponding tab title and then follow the steps to complete the task.
- To keep your changes, click Save.
- To discard your changes, click Cancel.

Result



Note:

If this is the first policy you are creating, you must click Done and designate the first policy as the default before saving on the CIBA Request Policies window. You can change the default as needed when you create additional policies.

OAuth attribute mapping using a datastore

Although an optional configuration, you can map OAuth attributes using a datastore in the PingFederate administrative console.

About this task

This optional configuration is the same for all OAuth grant mapping and token mapping configurations.

Steps

- 1. For grant mapping, go to Authentication # OAuth # IdP Adapter Grant Mapping. For more information, see Grant contract mapping on page 528.
- 2. For token mapping, go to Applications # OAuth # Access Token Mapping. For more information, see *Token mapping* on page 549.

OAuth client session management

When an organization opens web-based protected resources to their remote employees, business partners, and customers, it has limited control over the end-user devices. To minimize security risk, both the IT administrators and end users desire session management with tight controls.

PingFederate provides an Asynchronous Front-Channel Logout endpoint and a Back-Channel Session Revocation Web Service to help OAuth clients, such as PingAccess, to terminate sessions when end users log out and to prevent unauthorized access until the end users log in again.

PingAccess works out-of-the-box with PingFederate, taking full advantages of these two features.

Asynchronous Front-Channel Logout

Asynchronous Front-Channel Logout provides OAuth clients the capability to initiate single logout (SLO) requests to sign off associated SLO-enabled SAML 2.0 or WS-Federation sessions.

The Asynchronous Front-Channel Logout endpoint is /idp/startSLO.ping. Optionally, clients can add end-user sessions to a revocation list on logout and query the revocation list through the Back-Channel Session Revocation endpoint.



The Asynchronous Front-Channel Logout endpoint is also published in the OpenID Connect metadata at the /.well-known/openid-configuration endpoint. Look for ping end session endpoint in the metadata.

On a per-client basis, you can configure PingFederate to send logout requests, using the browser, to PingAccess and additional requests to other relying parties.

When you select the PingAccess option, PingFederate sends logout requests, using the browser, to the OpenID Connect logout endpoint on PingAccess (/pa/oidc/logout.png) to sign off other domains previously called by the session. For more information, see *OpenID Connect endpoints* in the PingAccess documentation.

In addition, when signing off an SLO-enabled SAML 2.0 or WS-Federation session, as the service provider (SP)-initiated logout request reaches the PingFederate identity provider (IdP) server, the same logout process applies as well. Depending on the enterprise architecture, this could further improve single sign-on (SSO) and logout use cases.

Back-Channel Session Revocation

Back-Channel Session Revocation allows OAuth clients, such as PingAccess, to query the revocation status of their sessions by sending HTTP GET requests to the session revocation endpoint on PingFederate at /pf-ws/rest/sessionMgmt/revokedSris.

To access the session revocation endpoint, a client must be granted access to the Session Revocation API. It must also authenticate with its client secret or client certificate and include in the request the session identifier, which can be obtained from the access token or the ID token.

Back-Channel Session Revocation also allows the clients to revoke sessions by sending HTTP POST requests to the same session revocation endpoint. This gives application developers the flexibility to revoke sessions based on the logic of their applications.

For each session added to the revocation list, PingFederate retains its revocation status for a configurable lifetime. Access control and authentication requirements to revoke sessions are identical to those to query for the revocation status.

OAuth token exchange

By configuring the OAuth authorization server to support OAuth token exchange, the authorization server can exchange a client's security token for another type of token.

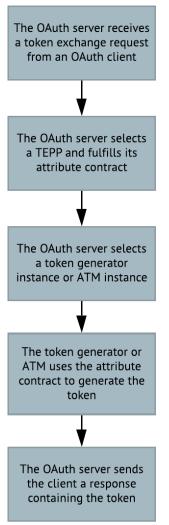
The OAuth token exchange allows resource servers to exchange access tokens for other security tokens that are required to call additional APIs, much like what the microservices architecture requires. PingFederate's native support of subject tokens and actor tokens opens new use cases around delegation and impersonation that enrich the end-user experience as resources flow through seamlessly among the back-end services used by the user-facing applications.

An OAuth token exchange transaction begins when an OAuth client sends the authorization server a request with the token exchange grant type. Then the authorization server gets the token exchange processor policy specified in the client's configuration.

The token exchange processor policy specifies which parameters from the token exchange request, and optionally which attributes from other sources, will be used. The policy always uses the subject token, but it can use an actor token too. The policy also specifies which token processor instance to use based on the request's subject token type and actor token type when present. During the transaction, the token processor instance transforms the subject token and optionally the actor token, parameters, and attributes into a token exchange processor policy attribute contract.

Depending on the type of token requested and what will consume it, the authorization server sends the attribute contract to a token generator instance or access token manager instance to generate the requested token.

This feature uses the protocol defined in the specification for *OAuth 2.0 Token Exchange*.



- The request contains the token exchange grant type parameter, a subject token, and a subject token type parameter.
- For delegation use cases, the request also contains an actor token and actor token type parameter. This diagram does not show that use case.
- The request may also contain resource, audience, scope, and requested token type parameters.
- If the client's configuration specifies a TEPP, then the OAuth server selects it. Otherwise, the OAuth server selects the default OAuth server TEPP.
- The OAuth server uses the TEPP's map of subject token types to select a token processor.
- The token processor uses the TEPP's attribute map to produce an attribute contract.
- If the request does not have an audience or resource parameter, then the OAuth server selects the default OAuth server ATM instance.
- If the request includes resource and/or audience parameters, then the OAuth server determines if they identify a single ATM or token generator instance.
- If the parameters do identify just one instance, then the OAuth server selects it; but if they do not identify just one, then the OAuth server returns an error.
- If the request does not include a requested token type parameter, then the default token type is generated.
- If the request includes a requested token type parameter, then that type of token is generated; but if the token generator or ATM does not support the requested token type, then the OAuth server returns an error.
- The response contains the requested token, and an issued token type parameter (the token's representation) and a token type parameter (the token's usage).

Notes:

- A token exchange processor policy (TEPP), includes a map of subject token types to token processor instances and a map of attributes from the request and other sources to a TEPP attribute contract.
- Optional resource parameters are used to identify an ATM instance or token generator instance.
- Optional audience parameters are used to identify another client that will use the new access token to get access to
 a resource. The OAuth server will select the ATM instance specified in that client's configuration.

Configuring OAuth token exchange

Configuring the OAuth authorization server to support OAuth token exchange involves configuring token exchange processor policies, token generator instances and token exchange generator groups, access token manager instances, and OAuth clients.

About this task

To configure OAuth token exchange, see the included topic links to perform the necessary steps.

Steps

1. Define token exchange processor policies to handle incoming token exchange requests. See *Defining token exchange processor policies* on page 596.

- 2. If you need token generator instances to generate the requested tokens, complete the following tasks.
 - a. Configure the token generator instances. See *Managing token generators* on page 975.
 - b. Create token exchange generator groups. See *Creating token exchange generator groups* on page 597.
 - c. Map the attributes from the token exchange processor policies to the attributes from the token generator instances. See *Mapping token exchange attributes to token generator attributes* on page 598.
- 3. Access token managers to generate the requested tokens.
 - a. Configure the access token manager instances. See *Managing access token management instances* on page 550.
 - b. Map the attributes from the token exchange processor policies to the attributes from the access token manager instances. See *Mapping token exchange attributes to access token manager attributes* on page 598.
- 4. Enable token exchange in the OAuth clients that will send the token exchange requests to the authorization server. See *Enabling token exchange in OAuth clients* on page 599.

Defining token exchange processor policies

To exchange security tokens, the OAuth authorization server needs at least one token exchange processor policy.

Before you begin

Before you define a token exchange processor policy, create the necessary token processor instances. See *Managing token processors* on page 953.

About this task

In the **Token Exchange Processor Policy Management** window, configure and define a token exchange processor policy.

Steps

- 1. Go to Applications # Token Exchange # Processor Polices to open the Token Exchange Processor Policy Management window.
- 2. Click Add Processor Policy.
 - Result: The Token Exchange Processor Policy window opens.
- 3. On the Manage Processor Policy tab, enter the policy ID and Name. Click Next.
 - Select the **Actor Token Required** check box if you want to specify whether the policy requires an actor token as well as a subject token in the token exchange requests from the clients.
- 4. On the Attribute Contract tab, add attributes to the attribute contract as needed. Click Next.

- 5. On the **Token Processor Mapping** tab, map a token processor to each subject token type or each combination of subject token type and actor token type:
 - a. Click the Map New Token Processor button.
 - Result: The **Token Processor Mapping** window opens.
 - b. On the **Token Types** tab, from the **Subject Token Processor** list, select the instance.
 - c. In the **Subject Token Type** field, enter the identifier.
 - d. If an actor token processor is required, from the Actor Token Processorlist, select the instance.
 - e. In the **Actor Token Type** field, enter the identifier. Click **Next**.
 - f. On the Attribute Sources & User Lookup tab, add additional attribute sources for contract fulfillment as needed. Click Next.
 - g. On the Contract Fulfillment tab, select the Source and Value for each attribute. Click Next.
 - h. On the **Issuance Criteria** tab, specify conditions that attributes must satisfy for PingFederate to exchange the token. Click Next.
 - i. On the **Summary** tab, review the token processor mapping. Click **Done**. Result: PingFederate returns you to the **Token Exchange Processor Policy** window.
- 6. On **Summary** tab, review the policy. Click **Done**.
 - Result: The Token Exchange Processor Policy Management window opens.
- 7. If you want to make the new token exchange processor policy the default policy, click **Set as Default** on the corresponding row in the table.
- 8. Click Save.

Creating token exchange generator groups

A token exchange generator group maps requested token types to your token generator instances. You can create multiple token exchange generator groups. If you assign resource URIs to the groups, clients can use the URI in the resource parameter of its requests to specify a group.

Before you begin

Before you create a token exchange generator group, configure the token generator instances. See Managing token generators on page 975.

About this task

In the **Generator Groups** window, create a generator group for a token exchange instance.

Steps

- 1. Go to the Applications # Token Exchange # Generator Groups window.
- 2. Click the Add Generator Group button.

The **Token Exchange Generator Group** window opens.

- 3. On the Manage Generator Group tab, enter the group ID, Name, and enter absolute Resource URIs.
- 4. On the Requested Token Type Mapping tab, from the Token Generator list, select an instance and enter the **Token Type**. Click **Add**.
- 5. Repeat step 4 for each type of token that you want the token exchange generator group to handle.



Important:

If this is the default token exchange generator group and clients will use the requested token type parameter to request specific types of tokens, then map all token types that clients can request.

6. On the **Summary** tab, review the token exchange generator group. Click **Done**.

The **Generator Groups** window opens.

- 7. If you want to make the new token exchange generator group the default group, click **Set as Default** in the **Action** column.
- 8. Click Save.

Mapping token exchange attributes to token generator attributes

When configuring the OAuth authorization server to exchange security tokens, if it uses token generator instances to create the requested tokens, then map the attributes in the attribute contract produced by the token exchange processor policy to the attributes created by the token generator instances.

Before you begin

Before you perform the following procedure:

- Define the token exchange processor policies. See Defining token exchange processor policies on page 596.
- Configure the token generator instances. See *Managing token generators* on page 975.

About this task

In the **Token Generator Mappings** window, map the attributes from a token exchange processor policy to the attributes from a token generator instance.

Steps

- 1. Go to Applications # Token Exchange # Token Generator Mappings to open the Token Generator Mappings window.
- 2. From the **Source Instance** list, select a token exchange processor policy.
- 3. From the **Target Instance** list, select a token generator from a token exchange generator group. Click **Add Mapping** button.
 - Result: The Mapping Configuration window opens.
- 4. On the **Attribute Sources & User Lookup** tab, add token generators additional attribute sources for contract fulfillment as needed. Click **Next**.
- 5. On the Token Contract Fulfillment tab, select a Source and Value for each attribute. Click Next.
- 6. On the **Issuance Criteria** tab, add and specify conditions that attributes must satisfy for PingFederate to exchange the token as needed. Click **Next**.
- 7. On the **Summary** tab, review the mapping configuration. Click **Done**.
 - Result: The **Token Generator Mappings** window opens.
- 8. Click Save.

Mapping token exchange attributes to access token manager attributes

When configuring the OAuth authorization server to exchange security tokens, if it uses an access token manager instances to generate requested tokens, then map the attributes in the attribute contract produced by the token exchange processor policy to the attributes in the tokens created by the access token manager instances.

Before you begin

Before you perform the following procedure:

- Define the token exchange processor policies. See Defining token exchange processor policies on page 596.
- Configure the access token managers instances. See Managing access token management instances on page 550.

About this task

In the **Access Token Mapping** window, map the attributes from a token exchange processor policy to the attributes from an access token manager instance.

Steps

- 1. Go to Applications # OAuth # Access Token Mapping.
- 2. In the Context section, from the Context list, select a token exchange processor policy.
- 3. From the **Access Token Manager** list, select an access token manager. Click **Add Mapping**. Result: The **Access Token Mapping** configuration window wizard opens.
- 4. On the **Attribute Sources & User Lookup**tab, add access token manager attribute sources for contract fulfillment as needed. Click **Next**.
- 5. On the Contract Fulfillment tab, select a Source and Value for each attribute. Click Next.
- 6. On the **Issuance Criteria** tab, add and specify conditions that attributes must satisfy for PingFederate to exchange the token as needed. Click **Next**.
- 7. On the **Summary** tab, review the access token mapping. Click **Done**.
 - Result: The Access Token Mapping window opens.
- 8. Click Save.

Enabling token exchange in OAuth clients

After configuring the OAuth authorization server to exchange tokens, enable token exchange on each OAuth client that will send the authorization server token exchange requests.

Before you begin

Before you perform the following procedure, define the token exchange processor policy that the OAuth server uses when it receives a token exchange request from the client. See *Defining token exchange* processor policies on page 596.

About this task

In the Client window, enable OAuth token exchange in an OAuth client.

Steps

- Go to Applications # OAuth # Clients to open the Clients window. Click the link of the client in the Client ID column.
 - Result: The **Client** window opens.
- 2. In the Allowed Grant Types section, select the Token Exchange check box.
- 3. In the **Token Exchange** section, from the **Processor Policy** list, select the token exchange processor policy.
- 4. Click Save.
 - Result: This will take you back to the Clients window.
- 5. Click Save.

Security management

The **Security** menu provides access to security and infrastructure-related settings. Depending on the setup of PingFederate, menu items vary.

Security management consists of the following sections:

- Certificate and key management on page 600
- System integration on page 631
- Account lockout protection on page 637
- Password spraying prevention on page 638
- Implementing a MasterKeyEncryptor using AWS KMS on page 639

Certificate and key management

PingFederate administrative users can manage keys and certificates for various purposes.

The PingFederate administrative console provides a suite of configuration wizards for administrators to manage keys and certificates. Tasks include:

- Managing trusted certificate authorities (CAs)
- Managing server certificates for the administrative port and runtime ports
- Managing client certificates for mutual TLS authentication
- Managing signing and decryption keys and certificates
- Managing OAuth and OpenID Connect keys
- Managing certificates from partners
- Configuring certificate revocation settings
- Managing partner metadata URLs
- Rotating system keys

You can configure PingFederate to use a hardware security module (HSM) for cryptographic material storage and operations. Standards such as the Federal Information Processing Standard (FIPS) 140-2 require the storage and processing of all keys and certificates on a certified cryptographic module.



Management of keys and certificates is restricted to administrative users with the Crypto Admin administrative role (see Administrative accounts on page 842).

See subsequent topics for configuration steps.

Manage trusted certificate authorities

Use the Trusted CAs window's functionality to import, export, review, and remove certificate authorities (CAs).

You can import your federation partner's CA certificate or self-signed certificates into PingFederate's global trust list on Security # Certificate & Key Management # Trusted CAs. If the CA is not one of the major authorities, you might also need to import the certificate from the CA that signed the partner certificate.



Note:

If a required CA certificate is already available from the Java runtime, you do not need to import the same certificate into the PingFederate store.

Importing trusted certificate authorities

Import your federation partner's certificate authority (CA) certificate or self-signed certificates into PingFederate's global trust list.

Steps

- 1. On the **Trusted CAs** window, click **Import**.
- 2. On the **Import Certificate** window, choose the applicable certificate file.



Note:

If PingFederate is integrated with a hardware security module (HSM) from Thales in hybrid mode, select the storage facility of the certificate from the Cryptographic Provider list.

- Select HSM to store the certificate in the HSM.
- Select Local Trust Store to store the certificate in the local trust store managed by PingFederate.

3. On the **Summary** window, review your configuration, amend as needed, and click **Save**.

Exporting trusted certificate authorities

Export your federation partner's certificate authority (CA) certificate or self-signed certificates as desired.

Steps

- 1. On the **Trusted CAs** window, select**Action** # **Export** for the certificate.
- 2. On the Export Certificate window, click Next.
- 3. On the Export & Summary window, click Export to save the certificate file and then click Done.

Reviewing trusted certificate authorities

Review certificates to ensure you've selected the correct ones.

Steps

- 1. On the **Trusted CAs** window, select the certificate by its serial number.
- 2. Review the selected certificate in the pop-up window.
- 3. When finished, close the pop-up window.

Removing trusted certificate authorities

Remove certificates from the Trusted CAs window when necessary.

Steps

1. On the **Trusted CAs** window, select **Action** # **Delete** for the certificate.



Note:

To cancel the removal request, select **Action # Undelete** for the certificate.

2. Click **Save** to confirm your action.

Manage SSL server certificates

Manage the certificates used to access the PingFederate administrative console and the runtime server.

Use Security # Certificate & Key Management # SSL Server Certificates to establish and maintain the certificates presented for access to the PingFederate administrative console or the administrative API and for incoming HTTPS connections at runtime.

The first system-generated certificate is the default certificate for both the administrative console and the runtime server. As multiple certificates are created, you can activate or deactivate them for the administrative console, the runtime server, or both. Additionally, you can select any of them as the new default certificate for the administrative console, the runtime server, or both at a later time.

When creating a certificate, you can add additional domain names through the use of the Subject Alternative Names field. Furthermore, if a user agent includes the host name that it intends to reach as part of the TLS handshake, PingFederate selects the applicable certificate based on the provided Server Name Indication (SNI) information. The selection looks at the common name and subject alternative names of each activated certificate. If PingFederate finds no match, it serves the default certificate. If PingFederate finds multiple matches, it serves the certificate with the better match.

SSL Server Certificates configuration

Certificate	Common name	Subject alternative names	Activation status
#1	www.example.com	(None)	Administrative console and runtime server

Certificate	Common name	Subject alternative names	Activation status
#2	www.example.org	*.example.org and test.example.local	Administrative console and runtime server
#3	www.example.info	*.example.info and *.example.com	Administrative console and runtime server
#4	admin.example.local	(None)	Administrative console (Default) and runtime server
#5	runtime.example.loca	1(None)	Administrative console and runtime server (Default)

Runtime behavior

Request type	Host name from SNI	Certificate served
Administrative or runtime	www.example.com	The host name from the SNI is an exact match to the common name of certificate #1 and a partial match to the second subject alternative name (*.example.org) of certificate #3.
		An exact match is a better match, so PingFederate serves certificate #1.
Administrative or runtime	www.example.org	The host name from the SNI is an exact match to the common name of certificate #2.
		PingFederate serves certificate #2.
Administrative or runtime	sso.example.org	The host name from the SNI is a partial match to the first subject alternative name (*.example.org) of certificate #2. There is no other exact or partial match.
		PingFederate serves certificate #2.
Administrative or runtime	sso.example.info	The host name from the SNI is a partial match to the first subject alternative name (*.example.info) of certificate #3. There is no other exact or partial match.
		PingFederate serves certificate #3.
Administrative or runtime	sso.example.com	The host name from the SNI is a partial match to the second subject alternative names (*.example.com) of certificate #3. There is no other exact or partial match.
		PingFederate serves certificate #3.
Administrative	www.example.local	The host name from the SNI does not match any configured certificate.
		PingFederate serves certificate #4, the default certificate for the administrative console.

Request type	Host name from SNI	Certificate served
Runtime	localhost	The host name from the SNI does not match any configured certificate.
		PingFederate serves certificate #5, the default certificate for the runtime server.



If PingFederate finds multiple certificates of the same matching quality, it returns one of them in the TLS handshake. This response should not impact the user agent because either the common name or one of the subject alternative names matches the host name of the request. If PingFederate should always serve a particular certificate for any given host name, ensure that the common name and any configured subject alternative names do not overlap among multiple certificates.

Creating a new certificate

Use the SSL Server Certificates window's functionality to generate customized certificates.

Steps

- 1. On the SSL Server Certificates window, click Create new.
- 2. On the **Create Certificate** window, enter the required information.

For information about each field, refer to the following table.

Field	Description
Common Name	The common name (CN) identifying the certificate.
Subject Alternative Names	The additional DNS names or IP addresses that can be associated with the certificate.
Organization	The organization (O) or company name creating the certificate.
Organizational Unit	The specific unit within the organization (OU).
City	The city or other primary location (L) where the company operates.
State	The state (ST) or other political unit encompassing the location.
Country	The country (C) where the company is based.
Validity (days)	The time during which the certificate is valid.
Cryptographic Provider	The storage facility of the certificate.
	Applicable and visible only when PingFederate is integrated with an HSM in hybrid mode.
	 Select HSM to store the certificate in the HSM.
	 Select Local Trust Store to store the certificate in the local trust store managed by PingFederate.
Key Algorithm	A cryptographic formula used to generate a key. PingFederate uses either of two algorithms, RSA or EC.
Key Size (bits)	The number of bits used in the key. (RSA-1024, 2048 and 4096; and EC-256, 384 and 521.)

Field	Description
Signature Algorithm	The signing algorithm of the certificate. (RSA-SHA256, SHA384, and SHA512; and ECDSA-SHA256, SHA384, and SHA512.)

- 3. When finished, click Next.
- 4. On the Summary window, review your configuration, amend as needed, and click Save.

Importing a certificate and its private key

Use the SSL Server Certificates window's functionality to import certificates.

Steps

- 1. On the SSL Server Certificates window, click Import.
- 2. On the **Import Certificate** window, choose the applicable certificate file and enter its password.



If PingFederate is integrated with a hardware security module (HSM) from Thales, you cannot use an elliptic curve (EC) certificate as an SSL server certificate. You must select a certificate that uses the RSA key algorithm.

- 3. If PingFederate is integrated with an HSM in hybrid mode, select the storage facility of the certificate from the Cryptographic Provider list.
 - Select HSM to store the certificate in the HSM.
 - b. Select Local Trust Store to store the certificate in the local trust store managed by PingFederate.
- 4. On the Summary window, review your configuration, amend as needed, and click Save.

Creating a certificate-authority signing request (CSR)

Use the SSL Server Certificates window's functionality to generate a CSR file for a certificate.

Steps

1. On the SSL Server Certificates window, select Action # Certificate Signing for the certificate.



Note:

This selection is inactive if you have not yet saved a newly created or imported certificate. Click Save and then return to this window to initiate the process.

The selection is also inactive if a previously signed certificate has been revoked. Because the revocation may indicate that the private key has been compromised, the best practice is to import or create a replacement certificate for certificate signing.

- 2. On the Certificate Signing window, select the Generate CSR option.
- 3. On the Generate CSR window, click Export to save the CSR file, and then click Done.



Note:

Once saved, you can submit this CSR file to a certificate authority (CA) for a CA-signed certificate.

Importing a certificate-authority response (CSR response)

Use the SSL Server Certificates window's functionality to import CSR response files for certificates.

Steps

1. On the SSL Server Certificates window, select Action # Certificate Signing for the certificate.

- 2. On the **Certificate Signing** window, select the **Import CSR Response** option.
- 3. On the **Import CSR Response** window, choose the applicable CSR response file.
- 4. On the **Summary** window, review your configuration, and click **Save**.

Exporting a certificate

Use the **SSL Server Certificates** window's functionality to export a certificate with or without its private key.

Steps

- 1. On the SSL Server Certificates window, select Action # Export for the certificate.
- 2. On the **Export Certificate** window, select the export type.
 - Select Certificate Only to export the selected certificate without its private key. This is the default choice.
 - b. Select **Certificate and Private Key** to export the selected certificate with its private key.



Note:

This export contains the private key of the certificate. You must also enter an encryption password.

If the selected certificate is stored in a hardware security module (HSM), the **Certificate and Private Key** option does not apply.

3. On the Export & Summary window, click Export to save the certificate file, and then click Done.

Reviewing a certificate

Use the **SSL Server Certificates** window's functionality to look over a particular certificate.

Steps

- 1. On the SSL Server Certificates window, select the certificate by its serial number.
- 2. Review the selected certificate in the pop-up window.
- 3. When finished, close the pop-up window.

Activating or deactivating a certificate

Use the **SSL Server Certificates** window's functionality to configure whether to activate or deactivate a certificate.

Steps

- On the SSL Server Certificates window, select the relevant option under Action for the certificate.
 Any certificate can be activated for the administrative console, the runtime server, or both. When multiple certificates are activated for the administrative console (or the runtime server), you can deactivate any of them as long as one certificate remains active. Additionally, you may select any of them as the default certificate.
- 2. Click **Save** to keep your configuration.

Use the SSL Server Certificates window's functionality to delete unwanted certificates.

Steps

1. On the SSL Server Certificates window, select Action # Delete for the certificate.



Note:

If the selected certificate is activated for the administrative port, the runtime port, or both, the **Delete** option does not apply.

To cancel the removal request, select **Action** # **Undelete** for the certificate.

2. Click **Save** to confirm your action.

Manage SSL client keys and certificates

On Security # Certificate & Key Management # SSL Client Keys & Certificates, you can create and manage your authentication private keys and the certificates your server presents as clients in an outbound SSL/TLS transaction.

The SSL Client Keys & Certificates window enables you to manage certificates and CSRs in multiple ways. The window's functionality allows you to create, import, export, review, and delete certificates, as well as create CSRs and import CSR responses.

Creating new certificates

Use the functionality found in the SSL Client Keys & Certificates window to create new, customized certificates.

Steps

- 1. On the SSL Client Keys & Certificates window, click Create new.
- 2. On the **Create Certificate** tab, enter the required information.

For information about each field, refer to the following table.

Field	Description
Common Name	The common name (CN) identifying the certificate.
Subject Alternative Names	The additional DNS names or IP addresses possibly associated with the certificate.
Organization	The organization (O) or company name creating the certificate.
Organizational Unit	The specific unit within the organization (OU).
City	The city or other primary location (L) where the company operates.
State	The state (ST) or other political unit encompassing the location.
Country	The country (C) where the company is based.
Validity (days)	The time during which the certificate is valid.
Key Algorithm	A cryptographic formula used to generate a key. PingFederate uses either of two algorithms, RSA or EC.
Key Size (bits)	The number of bits used in the key. (RSA-1024, 2048 and 4096; and EC-256, 384 and 521.)
Signature Algorithm	The signing algorithm of the certificate. (RSA and ECDSA-SHA256, SHA384, and SHA512.)

- 3. When finished, click Next.
- 4. On the **Summary** tab, review your configuration, amend as needed, and click **Done**.

Importing certificates and their private keys

Use the functionality of the SSL Client Keys & Certificates window to import your certificate key files into PingFederate.

Steps

- 1. On the SSL Client Keys & Certificates window, click Import.
- 2. On the **Import Certificate** tab, choose the applicable certificate file and enter its password.



If PingFederate is integrated with an HSM in hybrid mode, select the storage facility of the certificate from the Cryptographic Provider list.

- Select **HSM** to store the certificate in the HSM.
- Select Local Trust Store to store the certificate in the local trust store managed by PingFederate.
- 3. On the **Summary** window, review your configuration, amend as needed, and click **Done**.

Creating a certificate signing request (CSR)

Use the Certificate Signing functionality to generate and save a CSR file in order to submit it to a certificate authority (CA) for a signed certificate.

Steps

1. On the SSL Client Keys & Certificates window, select Certificate Signing for the certificate.



This selection is inactive if you have not yet saved a newly created or imported certificate. Click Save and then return to this window to initiate the process.

The selection is also inactive if a previously signed certificate is revoked. Because the revocation could indicate that the private key is compromised, the best practice is to import or create a replacement certificate for certificate signing.

- 2. On the **Certificate Signing** tab, select the **Generate CSR** option.
- 3. On the Generate CSR tab, click Export to save the CSR file, and then click Done.



Note:

Once saved, you can submit this CSR file to a certificate authority for a CA-signed certificate.

Importing a certificate-authority response (CSR response)

Use the Certificate Signing functionality to import your own CSR response file into PingFederate.

Steps

- 1. On the SSL Client Keys & Certificates window, select Certificate Signing for the certificate.
- 2. On the Certificate Signing tab, select the Import CSR Response option.
- 3. On the **Import CSR Response** tab, choose the applicable CSR response file.
- 4. On the **Summary** tab, review your configuration, and click **Save**.

Use the **Export** functionality to export a certificate and, if desired, its private key.

Steps

- 1. On the SSL Client Keys & Certificates window, select Export for the certificate.
- 2. On the **Export Certificate** tab, select the export type.
 - Select Certificate Only to export the selected certificate without its private key. This is the default
 - Select Certificate and Private Key to export the selected certificate with its private key.



CAUTION:

This export contains the private key of the certificate. You must also enter an encryption password.

If the selected certificate is stored in a hardware security module (HSM), the Certificate and Private Key option does not apply.

3. On the Export & Summary window, click Export to save the certificate file, and then click Done.

Reviewing certificates

Take a closer look at individual certificates to ensure their properties match your needs.

Steps

- 1. On the SSL Client Keys & Certificates window, select the certificate by its serial number.
- 2. Review the selected certificate in the pop-up window.
- 3. When finished, close the pop-up window.

Removing certificates

Delete certificates you no longer need.

Steps

1. On the SSL Client Keys & Certificates window, select Delete for the certificate.



Note:

To cancel the removal request, select **Undelete** for the certificate.

2. Click **Save** to confirm your action.

Manage digital signing certificates and decryption keys

On Security # Certificate & Key Management # Signing & Decryption Keys & Certificates, you can create and maintain certificates and their respective key pairs for the purpose of signing outgoing requests, responses, assertions, and access tokens, and for the purpose of decryption.

Use separate certificates for signing and decryption.

After creating your certificates, if they remain as self-signed certificates, you can enable automatic certificate rotation. See Certificate rotation on page 609.

The optional automatic certificate rotation feature of PingFederate greatly reduces the cost of managing self-signed certificates.

PingFederate supports automatic certificate rotation for self-signed certificates created for signing SAML requests, responses, and assertions, or XML decryption for browser SSO and WS-Trust STS transactions on a per-certificate basis.



Certificate rotation is only available to self-signed certificates.

Certificate rotation happens over two stages, identified by the Creation Buffer and Activation Buffer settings.

- The Creation Buffer is the number of days ahead of expiry that PingFederate creates a new key pair and a new certificate.
- The Activation Buffer is the number of days ahead of expiry that PingFederate activates the certificate.

When you enable certificate rotation on a certificate, you can customize the values of the Creation Buffer and **Activation Buffer** settings. Alternatively, you can keep their default values, which are 25% and 10% of the original lifetime of the current certificate, respectively. The following examples illustrate the default values for both buffers based on a 100-day certificate and a 365-day certificate.

Current certificate	The default value for the Creation Buffer field	The default value for the Activation Buffer field	The rotation window
Self-signed certificate #1, valid for 100 days from January 1, 2017 to April 9, 2017	25 days ahead of expiry, which is March 16	10 days ahead of expiry, which is March 31	15 days from March 16 through March 30
Self-signed certificate #2, valid for 365 days from January 1, 2017 to December 31, 2017	91 days ahead of expiry, which is October 2	36 days ahead of expiry, which is November 26	55 days from October 2 through November 25

If the PingFederate server is shut down when the Creation Buffer threshold is reached for a given certificate, a new key pair and a new certificate are created if PingFederate is restarted during the rotation window.

In a clustered PingFederate environment, when the new signing certificate is ready, the administrative console displays a message to remind the administrators to replicate the new certificate to the engine nodes in System # Server # Cluster Management.

Although optional, you can turn on notifications for certificate events in **System # Monitoring &** Notifications # Runtime Notifications. When configured, PingFederate notifies the configured recipient when a new certificate is available and when it is activated. Depending on the role of the certificate, you can update your partner accordingly.

Connection and federation metadata

Certificate rotation uses a number of inherent capabilities which enable it to deploy new certificates to replace current certificates in enabled connections..

Certification rotation is a per-certificate configuration. When certificate rotation is enabled for a certificate and a new certificate using new key pairs becomes available, PingFederate deploys the new certificate to all enabled connections that use the original certificate. The actions taken by PingFederate vary depending on the role of the certificate.

Notifications

Although optional, you can turn on notifications for certificate events in **System # Monitoring &** Notifications # Runtime Notifications. When configured, PingFederate notifies the configured recipient when a new certificate is available and when it is activated. Depending on the role of the certificate, you can update your partner accordingly.

Signing certificate

When the Creation Buffer threshold is reached, a new certificate is created. For all web browser single sign-on (SSO) (SAML and WS-Federtion) connections using the same signing certificate, PingFederate starts including the new certificate (along with the current certificate) in their metadata. PingFederate keeps using the current certificate for signing until the remaining lifetime of the current certificate reaches the Activation Buffer threshold, at which point PingFederate starts signing with the new certificate and removes the previous certificate from the metadata.



Important:

To prevent SSO outages, partners must update their connections to use the new certificate to verify digital signatures before the **Activation Buffer** threshold is reached.

XML decryption

When a new certificate is made available, PingFederate performs the following tasks for all SAML 2.0 connections using the same decryption key:

- Pushes the current decryption key from primary to secondary
- Places the new certificate as the primary decryption key
- Updates the decryption key with the new certificate in the metadata
- Starts using the new decryption key to decrypt inbound messages. If the primary decryption key fails, PingFederate fails over to the secondary decryption key

When the remaining lifetime of the current certificate reaches the Activation Buffer threshold, the secondary decryption key is removed from the SAML 2.0 connections.

When PingFederate is configured to generate notifications for certificate events, PingFederate also notifies the configured recipient when the existing RSA decryption key is about to expire.



Important:

For XML decryption keys, PingFederate only supports the RSA key algorithm. When EC (elliptic curve) is selected as the **Key Algorithm** value on the **Certificate Rotation** tab, PingFederate does not update the SAML 2.0 connections and their metadata.



Important:

To prevent SSO outages, partners must update their connections to use the new certificate to encrypt messages before the **Activation Buffer** threshold is reached.

Federation metadata for Browser SSO connections

PingFederate updates the metadata for the applicable web browser SSO connections as soon as a new certificate is available.

To ensure that your partners are aware of the new certificate, you can provide their respective federation metadata by URLs or exports.

PingFederate runtime engine provides an endpoint (/pf/federation metadata.ping) to return metadata for web browser SSO connections. A service provider (SP) or an identity provider (IdP) is identified by its entity IDs using the PartnerSpId guery parameter or the PartnerIdpId guery parameter, respectively, as illustrated in the following examples.

Partner	Federation metadata URL to be given to the partner
An SP partner with an entity ID of SP1.	https://www.example.com:9031/pf/ federation_metadata.ping?PartnerSpId=SP1
An IdP partner with an entity ID of IdP1.	https://www.example.com:9031/pf/ federation_metadata.ping?PartnerIdpId=IdP1



The base URL for the PingFederate runtime engine is https://www.example.com:9031



Important:

In a clustered environment, because the console node is responsible for creating and applying the new certificates to all applicable connections, you must replicate the new certificate to the engine nodes in System # Server # Cluster Management when the new certificate is available, so that the federation metadata for these connections is updated accordingly.

The administrative console reminds you to replicate configuration when it detects configuration changes.

Metadata by manual export

Alternatively, you can export a metadata file for a connection from the Connections Management window or System # Protocol Metadata # Metadata Export.



Note:

PingFederate does not deploy new certificates or update metadata for inactive connections.

WS-Trust STS connections

For connections with only the WS-Trust security token service (STS) profile, you must export the new pending certificate and pass it to your partners out-of-band before the Activation Buffer threshold is reached.

If a connection contains both the Browser SSO and the WS-Trust STS profiles, the new certificate is included in the federation metadata for the Web Browser SSO profile. Your partner can reuse the certificate from the metadata by URL or manual export and apply it to its STS configuration.

Managing certificate rotation settings

Use the Signing & Decryption Keys & Certificates window to customize certificate rotation settings for your certificates.

About this task

Manage certificate rotation settings for self-signed certificates on Security # Certificate & Key Management # Signing & Decryption Keys & Certificates.

Steps

1. On the Signing & Decryption Keys & Certificates window, select Certificate Rotation for the applicable certificate.



Note:

Certificate rotation is only available to self-signed certificates.

- 2. Select the check box to turn on certificate rotation for the selected certificate, then click Next. If you want to turn off certificate rotation for the selected certificate, clear the check box and then click Save.
- 3. Optional: On the **Certificate Rotation** tab, modify the default values.

Field	Description
Creation buffer	The number of days ahead of expiry that PingFederate creates a new key pair and a new certificate.
	The default value is 25% of the original lifetime of the current certificate.
Activation buffer	The number of days ahead of expiry that PingFederate activates the certificate.
	The default value is 10% of the original lifetime of the current certificate.
Validity	The time during which the certificate is valid.
	The default value matches that of the current certificate.
Key Algorithm	A cryptographic formula used to generate a key. PingFederate uses either of two algorithms, RSA or EC.
	The default value matches that of the current certificate.
	Important: For XML decryption keys, PingFederate only supports the RSA key algorithm. When EC (elliptic curve) is selected as the Key Algorithm value on the Certificate Rotation tab, PingFederate does not update the SAML 2.0 connections and their metadata.
Key Size	The number of bits used in the key. (RSA-1024, 2048 and 4096; and EC-256, 384 and 521.)
	The default value matches that of the current certificate.

Field	Description
Signature Algorithm	The signing algorithm of the certificate. (RSA and ECDSA-SHA256, SHA384 and SHA512.)
	The default value matches that of the current certificate.

4. On the Certificate Rotation Summary tab, review the rotation settings. Adjust as needed, and then click Save to turn on automatic certificate rotation for this certificate.

Managed SP connection to PingOne for Enterprise and signing certificate

Use managed service provider (SP) connections to PingOne for Enterprise to automatically rotate signing certificates being used by it.

PingFederate automatically rotates the signing certificate used by the managed SP connection to PingOne for Enterprise.



A managed SP connection to PingOne for Enterprise is a connection created either as part of the initial setup or the System # External Systems # Connect to PingOne for Enterprise configuration wizard in PingFederate 8.0 or later.

The certificate rotation settings are as follows.

Field	Values
Creation Buffer (days)	90
Activation Buffer (days)	30
Validity (days)	1095
Key Algorithm	RSA
Key Size	2048
Signature Algorithm	RSA SHA256

If the signing certificate should be manually rotated instead, disable automatic certificate rotation. See Managing certificate rotation settings on page 612.



Note:

After making changes, the administrative console prompts for confirmation whether to update PingOne for Enterprise or to disconnect from PingOne for Enterprise in a banner message. See Managing PingOne for Enterprise settings on page 921.

Creating new certificates

Use the functionality found in the Signing & Decryption Keys & Certificates window to create new, customized certificates.

Steps

1. On the Signing & Decryption Keys & Certificates window, click Create new.

For information about each field, refer to the following table.

Field	Description
Common Name	The common name (CN) identifying the certificate.
Subject Alternative Names	The additional DNS names or IP addresses possibly associated with the certificate.
Organization	The organization (O) or company name creating the certificate.
Organizational Unit	The specific unit within the organization (OU).
City	The city or other primary location (L) where the company operates.
State	The state (ST) or other political unit encompassing the location.
Country	The country (C) where the company is based.
Validity (days)	The time during which the certificate is valid.
Key Algorithm	A cryptographic formula used to generate a key. PingFederate uses either of two algorithms, RSA or EC.
Key Size (bits)	The number of bits used in the key. (RSA-1024, 2048 and 4096; and EC-256, 384 and 521.)
Signature Algorithm	The signing algorithm of the certificate. (RSA and ECDSA-SHA256, SHA384, and SHA512.)

- 3. When finished, click **Next**.
- 4. On the **Summary** window, review your configuration, amend as needed, and click **Done**.

Importing certificates and their private keys

Use the functionality of the **Signing & Decryption Keys & Certificates** window to import your certificate key files into PingFederate.

Steps

- 1. On the Signing & Decryption Keys & Certificates window, click Import.
- 2. On the **Import Certificate** tab, choose the applicable certificate file and enter its password.



Note:

If PingFederate is integrated with an HSM in hybrid mode, select the storage facility of the certificate from the **Cryptographic Provider** list.

- Select **HSM** to store the certificate in the HSM.
- Select Local Trust Store to store the certificate in the local trust store managed by PingFederate.
- 3. On the **Summary** window, review your configuration, amend as needed, and click **Done**.

Creating a certificate signing request (CSR)

Use the Certificate Signing functionality to generate and save a CSR file in order to submit it to a certificate authority (CA) for a signed certificate.

Steps

1. On the Signing & Decryption Keys & Certificates window, select Certificate Signing for the certificate.



Note:

This selection is inactive if you have not yet saved a newly created or imported certificate. Click Save and then return to this window to initiate the process.

The selection is also inactive if a previously signed certificate is revoked. Because the revocation could indicate that the private key is compromised, the best practice is to import or create a replacement certificate for certificate signing.

- 2. On the **Certificate Signing** tab, select the **Generate CSR** option.
- 3. On the Generate CSR tab, click Export to save the CSR file, and then click Done.



Note:

Once saved, you can submit this CSR file to a certificate authority for a CA-signed certificate.

Importing a certificate-authority response (CSR response)

Use the Certificate Signing functionality to import your own CSR response file into PingFederate.

Steps

- 1. On the Signing & Decryption Keys & Certificates window, select Certificate Signing for the certificate.
- 2. On the Certificate Signing tab, select the Import CSR Response option.
- 3. On the **Import CSR Response** tab, choose the applicable CSR response file.
- 4. On the **Summary** tab, review your configuration, and click **Save**.

Exporting certificates

Use the **Export** functionality to export a certificate and, if desired, its private key.

Steps

- 1. On the Signing & Decryption Keys & Certificates window, select Export for the certificate.
- 2. On the **Export Certificate** tab, select the export type.
 - Select Certificate Only to export the selected certificate without its private key. This is the default choice.
 - Select Certificate and Private Key to export the selected certificate with its private key.



CAUTION:

This export contains the private key of the certificate. You must also enter an encryption password.

If the selected certificate is stored in a hardware security module (HSM), the Certificate and **Private Key** option does not apply.

3. On the Export & Summary window, click Export to save the certificate file, and then click Done.

Take a closer look at individual certificates to ensure their properties match your needs.

Steps

- 1. On the **Signing & Decryption Keys & Certificates** window, select the certificate by its serial number.
- 2. Review the selected certificate in the pop-up window.
- 3. When finished, close the pop-up window.

Reviewing a certificate's usage

Take a look at a certificate's usage data to get a sense of how often it's used.

Steps

1. On the Signing & Decryption Keys & Certificates window, select Check Usage for the certificate.



Note:

If the certificate is not used by any configuration, the **Check Usage** option does not apply.

- 2. Review the information in the pop-up window.
- 3. When finished, close the pop-up window.

Removing certificates

Delete certificates you no longer need.

Steps

1. On the Signing & Decryption Keys & Certificates window, select Delete for the certificate.



Note:

To cancel the removal request, select **Undelete** for the certificate.

2. Click **Save** to confirm your action.

Keys for OAuth and OpenID Connect

You can use keys to manage a number of security roles in PingFederate.

On Security # Certificate & Key Management # OAuth & OpenID Connect Keys, you can specify whether PingFederate should use static or dynamically rotating keys for OAuth and OpenID Connect.



Note:

When using dynamically rotating keys, the number of key sets in memory is set to three for both signing and encryption keys. This number is not configurable. The key sets include pending, active, and retired. At each rotation cycle, a new set of pending keys is generated. The original pending set becomes the active set, the active set becomes the retired set, and the old retired set goes away. All three sets are published for signing keys. For encryption keys, only the active key set is published. The rotation period and RSA key size are configurable in the file <pf install>/pingfederate/server/default/data/configstore/jwks-endpoint-configuration.xml.

The keys are used in the following manner.

PingFederate role	Key usages
Authorization Server (AS)	Sign self-contained access tokens for relying parties (RPs).

PingFederate role	Key usages
OpenID Provider (OP)	Sign ID tokens for RPs.
Relying Party (RP)	Sign JSON web tokens (JWTs) for authentication, sign OpenID Connect request objects, decrypt ID tokens, or any combination.

Configuring static signing keys

Determine when to use static and dynamically rotating keys in order to sign tokens as needed.

About this task

Specify whether PingFederate should use static or dynamically rotating keys to sign self-contained access tokens, ID tokens, JSON web tokens (JWTs) for client authentication, and JWTs for OpenID Connect request objects.

Steps

- 1. Go to Security # Certificate & Key Management # OAuth & OpenID Connect Keys.
- 2. Select the Enable Static Keys check box to use static keys for OAuth and OpenID Connect.



Clear this check box to let PingFederate generate and rotate keys automatically for OAuth and OpenID Connect.

The Enable Static Keys check box is not selected by default.

Result:

Once selected, the administrative console displays the following fields under the "Signing Keys" headingr.

Key Type	Active	Previous	Publish Certificate
EC with P-256 curve	Optional	Optional	Optional
EC with P-384 curve	Optional	Optional	Optional
EC with P-521 curve	Optional	Optional	Optional
RSA	Required	Optional	Optional

- 3. Follow these steps to complete the configuration under Signing Keys.
 - a. For the RSA key type, select an active signing key and optionally a previous signing key.



Note:

If you don't find the desired signing key, click Manage Certificates to create it.

There is no default selection.

Result:

The active signing key and the previous signing key (if configured) are published at the PingFederate JSON Web Key (JWK) Set endpoint /pf/JWKS.

b. For each applicable EC (elliptic curve) key type, select an active signing key and optionally a previous signing key.



Note:

If you don't find the desired signing key, click **Manage Certificates** to create it. Alternatively, complete the configuration, create the desired signing keys later, and then update the configuration afterward.

There is no default selection.

Result:

The active signing key and the previous signing key (if configured) are published at the PingFederate JWKS endpoint /pf/JWKS.

c. Optional: For any key type for which you have selected an active signing key (with or without a previous signing key), select the **Publish Certificate** check box to publish the certificates associated with the active signing key and the previous signing key (if configured) at the PingFederate JWKS endpoint /pf/JWKS.



Tip:

For each applicable signing key, its associated chain of certificates is published as the x5c parameter value.

The **Publish Certificate** check boxes are not selected by default.

4. Click Save.

Result



Important:

When static keys are enabled, PingFederate uses only static signing keys to sign ID tokens for OAuth clients or to sign JWTs for authentication or request objects (or both) for authorization servers; dynamic keys are not used and are not returned by the PingFederate JWKS endpoint /pf/JWKS. Signing algorithms associated with EC key types not configured with an active static signing key are hidden.

For existing clients and identity provider (IdP) connections, if you have previously selected a certain signing algorithm associated with an EC key type (for example, **ECDSA using P256 Curve and SHA-256**) without enabling static keys and then subsequently decide to enable static keys without selecting an active signing key for such EC key type (**EC with P-256 curve** in this example), transactions that involve that signing

algorithm will fail. When you revisit the configuration, the administrative console displays an error message. Your options are as follows:

OAuth clients

- Click Save to update the value of the ID Token Signing Algorithm setting to Default, which is the equivalent of selecting RSA using SHA-256 from the list.
- Select a different value from the **ID Token Signing Algorithm** list and save the configuration.
- Ignore the error and click Cancel without updating the configuration. Note that runtime errors persist until the configuration issue is resolved.

These options are applicable to individual clients on the **Client** window and the default setting configured for all clients created via the Dynamic Client Registration protocol on the Client Configuration Defaults window.

OpenID Connect IdP connections

- Select a different value from the Authentication Signing Algorithm list or the Request Signing Algorithm list (or both) and save the configuration.
- Ignore the error and click Cancel without updating the configuration. Runtime errors persist until the configuration issue is resolved.

These options are applicable to individual OpenID Connect IdP connections on the OpenID Provider Info window.

Configuring static decryption keys

You can specify whether PingFederate should use static or dynamically rotating keys to decrypt asymmetrically-encrypted ID tokens...

Steps

- 1. Go to Security # Certificate & Key Management # OAuth & OpenID Connect Keys.
- 2. Select the **Enable Static Keys** check box to use static keys for OAuth and OpenID Connect.



Clear this check box to let PingFederate generate and rotate keys automatically for OAuth and OpenID Connect. The Enable Static Keys check box is not selected by default.

Result:

Once selected, the administrative console displays the following fields under "Decryption Keys".

Key Type	Active	Previous	Publish Certificate
EC with P-256 curve	Optional	Optional	Optional
EC with P-384 curve	Optional	Optional	Optional
EC with P-521 curve	Optional	Optional	Optional
RSA	Optional	Optional	Optional

- 3. Follow these steps to configure "Decryption Keys".
 - a. For each applicable key type, select an active decryption key and optionally a previous decryption key.



Note:

If the desired decryption key is not found, click **Manage Certificates** to create it. Alternatively, complete the configuration, create the desired decryption keys later, and then update the configuration afterward. There is no default selection.

Result:

The active decryption key is published at the PingFederate JSON Web Key Set (JWKS) endpoint /pf/JWKS.

b. Optional: For any key type for which you have selected an active decryption key (with or without a previous decryption key), select the **Publish Certificate** check box to publish the certificates associated with the active decryption key at the PingFederate JWKS endpoint /pf/JWKS.



Tip:

Each applicable decryption key's associated chain of certificates is published as the x5c parameter value.

The **Publish Certificate** check boxes are not selected by default.



When static keys are enabled, you must also select an active signing key for the RSA key type.

4. Under "Signing Keys", select an active key for the RSA key type.



Note:

If the desired key is not found, click **Manage Certificates** to create it. There is no default selection.

Result:

The active signing key is published at the PingFederate JWKS endpoint /pf/JWKS.

5. Click Save.

Result



Important:

When static keys are enabled, PingFederate uses only static decryption keys to decrypt asymmetricallyencrypted ID tokens it receives from OpenID providers. Dynamic keys are not used and are not returned by the PingFederate JWKS endpoint /pf/JWKS.

The following snippet illustrates a sample response returned by the PingFederate JWKS endpoint when dynamic keys are used.

```
$ curl -s https://localhost:8031/pf/JWKS |python -m json.tool
  "keys": [
```

```
"kty": "EC",
      "kid": "I-ZbqeLPG2O5qxSf3n8yKmcGbWI",
      "use": "enc",
      "x": "AUSx-2vdfCjU90KohVs1peISnNUeDmGo3m0 x42PucBr-Gd-
mHKXQ8EjTeYgLhFB5SYMV5tntKiezayWkUt9Dodc",
      "y": "AIE6vQYcKdOfyQYzENYQ86MIAwSUo4GR_-
dn7m2MvRReXkotWOsFT1WKXi KjamqJIV2AwAUZL-IQj5mew451STM",
      "crv": "P-521"
    },
      "kty": "EC",
      "kid": "S2BbNNK9PtG0nA-EhU5BGpZ-OG8",
      "use": "enc",
      "x": "IKXASh9aDPJ1YaeXUww1YZnZ3kum WLKvZe8xiNW6W8",
      "y": "7 zp2AuY8MY4WEuneHEzV0cqW0buqcmMGVzRANQ0r2I",
      "crv": "P-256"
    },
      "kty": "EC",
      "kid": "t4-jKfmhEHn3mRc-080h3WKA2zE",
      "use": "enc",
      "x": "RiQkv ArGS7Zc8XsXp0VQpEWz9ZUlbLUWA0VbTcUjWIbOByceGhg-
tAj6dlFiorq",
     "y": "aHPQlrJPscdcuHtHokyr-70yBo4nUK-
BjWrJgisDxnKJQFLP6YK dfuOpuVYhFJ5",
      "crv": "P-384"
    },
      "kty": "RSA",
      "kid": "tVP7otNKgIWYep8LPBR3wD3tPNE",
      "use": "enc",
      "n":
 "hvHfiamhV4wGC9JHppJZjdKG5K3MvhWwo6PBsSQowGOTeILAbzO8Jfmp7nRxuujTE6k83RXNeWUvTwamGqShX
vjoNZD8Cv0Y9C3R4Ckj6dBL70Osk NfBR7MYmRA6dV0PJ5k4Lt vQveXMkylD9XuLFP-
gqooMXkB6FCCLqZZAi0voi3WQ7ECzSta3ke9F5VF17-4zVjRtJHjM9gGEhd5OkaZioqs9xBHeOrwhPbiPTsIA7v6
dpl dKixFTdQYIBMmIWGUyuB43XYq106z9CWoOcw",
     "e": "AQAB"
    },
    . . .
  1
```

When static keys are used, the PingFederate JWKS endpoint /pf/JWKS returns only the configured active keys. The following snippet illustrates a sample response returned by the PingFederate JWKS endpoint when an active key was selected for the **EC with P-384 curve** and **EC with P-521** curve key types.

```
"kid": "pE1XwX8Z6QYhAC7mjZ0OCn4DXAk",
      "use": "enc",
      "x": "ATCOsxg6ce437qMVlrqCyHPDE76hC0wP7Wwb7V8heai60LIDDvIJt-
evxTOGn7Iolo9PYET8-Bjhu5Zg5MNxOkF-",
 "AdvUA2YD2kn7COLkFIG2vL2k34CMv7VPxsvbgOJBL2exSziMGPw6YJp2eafuHlBom7bkjv3iFy5dTuGB7B28Z
     "crv": "P-521"
    },
 ]
}
```

Managing certificates from partners

Manage certificates for various connections involving signature verification, encryption, and back-channel authentication to effectively process messages to and from partners.

About this task

You receive certificates from partners for signature verification, encryption, and back-channel authentication. They are managed within connections.



Depending on the use cases, your connection to the partner might not require signature verification, encryption, inbound (SOAP) back-channel authentication by client certificate, or any such combinations. If so, the **Activation & Summary** window does not display the related administrative window.

Signature verification

Specify one or more certificates that PingFederate can use to validate the digital signatures found in inbound messages from your partners.

Steps

- 1. Select the connection to reach its **Activation & Summary** window.
- 2. Select Signature Verification Certificate.
- 3. Click Manage Certificates.



You can import, export, review, activate, deactivate, and remove certificates for signature verification on the Certificate Management window.

Encryption

Specify a certificate that PingFederate uses to encrypt outbound messages before delivering them to your partners.

Steps

- 1. Select the connection to reach its **Activation & Summary** window.
- 2. Select Select XML Encryption Certificate.
- 3. Click Manage Certificates.



Note:

You can import, export, review, activate, and remove certificates for encryption on the Certificate Management window.

Back-channel authentication

Specify a certificate that Pingfederate uses to authenticate inbound (SOAP) messages from your partners by their client certificates.

Steps

- 1. Select the connection to reach its **Activation & Summary** window.
- 2. Select SSL Verification Certificate.
- 3. Click Manage Certificates.



You can import, export, review, activate, and remove certificates for back-channel authentication on the Certificate Management window.

Configuring certificate revocation

Choose whether to utilize certificate revocation list (CRL) checking or Online Certificate Status Protocol (OCSP) checking as your preferred verification method.

About this task

By default at runtime, whenever a certificate revocation list (CRL) distribution-point URL is included within the certificate, PingFederate attempts to retrieve a CRL to verify that a signing certificate is not revoked. Optionally, on Security # Certificate & Key Management # Certificate Revocation Checking, you can enable and configure OCSP checking as the preferred verification method, depending on your requirements. For more information, see Certificate validation on page 87.

You can use OCSP in place of CRL checking, or retain CRLs as a backup method for failover.



When OCSP is enabled, CRL checking is not done independently, but only as a failover option for one or more OCSP failure conditions.

Steps

1. Optional: Configure OCSP.

For more information about each field, see the following table.

Field	Description	
Enable OCSP	Turns on OCSP certificate-revocation checking.	
	OCSP checking is not enabled by default.	
Default OCSP Responder URL	The location of a URL to use for certificate-revocation checking, a backup used only if the OCSP Responder URL is not contained in the certificate.	
Default OCSP Responder Signature Verification Certificate	Certificate used to verify that the returned certificate status was sent from the Default OCSP Responder—required if the certificate is not included in the response.	
	Click Manage Certificates to import the verification certificate.	

Field	Description
Do NOT allow Responder to use cached responses	When not selected, the OCSP Responder uses cached responses when available for the subject certificate for an indicated period of time—see the description for Next Update Grace Period .
	If checked, PingFederate sends a nonce in the request to the Responder, effectively requiring that the status of the certificate be determined in real time. This option is intended to enhance the prevention of Internet replay attacks (in addition to timestamping), where required.
	Important:
	Making this selection might slow down OCSP response time for a request and will increase general processing overhead at the Responder site.
	This check box is not selected by default.
This Update Grace Period (min)	To consider the response valid, the PingFederate server-clock time must correspond to the <thisupdate> timestamp in the OCSP response, plus or minus the number of minutes set for this field to compensate for clock variances.</thisupdate>
	The default value is 5 minutes.
Next Update Grace Period (min)	If the response includes a <nextupdate> timestamp indicating when updated certificate statuses are available, then PingFederate checks to ensure that the timestamp is not earlier than the current server time, adding this grace period to compensate for clock variances.</nextupdate>
	The default value is 5 minutes.
Responder Timeout (sec)	The allowable response time before the OCSP Responder URL is considered unavailable and processing continues. See the OCSP Responder is Unavailable setting.
	The default value is 5 seconds.
Response Caching	The number of hours that PingFederate caches the OCSP response.
Interval (hrs)	The default value is 48 hours.
Certificate is Unknown	The certificate does not fall under the purview of the certificate authority (CA) associated with the OCSP Responder. The choices indicate whether an unknown certificate is considered valid, or whether to try CRL checking.
	The default selection is Treat as Revoked .
OCSP Responder is Unavailable	Indicates what action to take if you cannot reach the Responder. The choices indicate whether an unknown certificate is considered valid, or whether to try CRL checking.
	The default selection is Treat as Valid .
OCSP Responder Returns Error	Indicates what action to take if the Responder returns an error. The choices indicate whether an unknown certificate is considered valid, or whether to try CRL checking.
	The default selection is Treat as Revoked .

For more information about each field, see the following table.

Field/Selection	Description
Enable CRL Checking	Enables CRL revocation checking.
	Note:
	CRL checking must remain enabled if any selections for OCSP Error Handling include failover. If OCSP is enabled and no CRL failover is specified, then this selection has no effect.
,	CRL revocation checking is enabled by default.
Treat Unretrievable CRLs as Revoked	If selected, PingFederate immediately aborts the processing associated with the certificate.
	If not selected, the server treats the certificate as valid but continues trying to retrieve the CRL.
	This check box is not selected by default.
Next Retry on Resolution Failure (min)	Specifies the number of minutes the server waits before trying to retrieve a CRL when the previous attempt failed—applies only when Treat Unretrievable CRLs as Revoked is unchecked.
	The default value is 1440 minutes, which is 24 hours.
Next Retry on Next Update Expiration (min)	How long the server waits before requesting a new CRL when the most recently retrieved CRL (in cache) has a next-update time in the past.
	Note:
	Certain actions in the administrative console, such as saving changes to an identity provider (IdP) adapter instance, reset the CRL cache. When this happens, PingFederate requests new CRLs for subsequent transactions as needed.
,	The default value is 60 minutes.
Verify CRL Signature	When selected (recommended), PingFederate verifies the CRL signature using the public key of the issuer, which must be in the certificate chain or in the list of Trusted CAs. See <i>Manage trusted certificate authorities</i> on page 600.
	This check box is selected by default.
Proxy Settings	If CRL checking is routed through a proxy server, specify the server's host DNS name or IP address and the port number. The same proxy information applies to OCSP checking, when enabled.

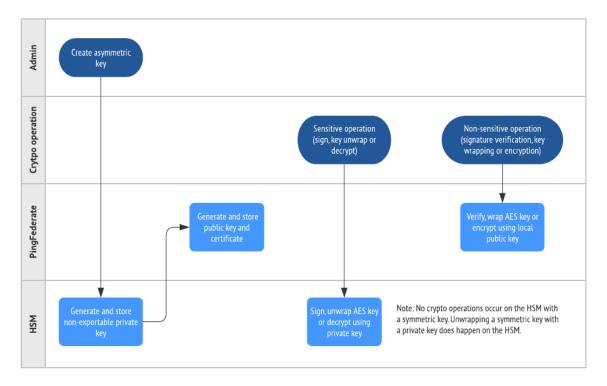
Use the PingFederate administrator functionality to determine whether to store keys and certificates on a hardware security module (HSM) or a local trust store.

About this task

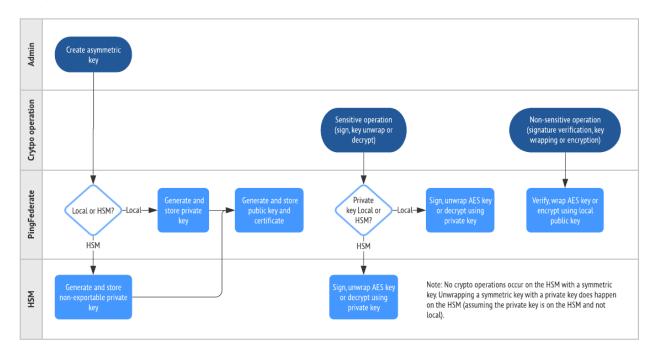
Administrators can enable the HSM hybrid mode, which provides the choice to store each relevant key and certificate on an HSM or the PingFederate-managed local trust store. This capability allows organizations to transition the storage of keys and certificates to a supported HSM to meet security requirements without the need to deploy a new PingFederate environment and mirror the setup.

The following images illustrate some general interactions between PingFederate and an HSM. Those interactions depend on whether you configure the HSM in hybrid mode.

PingFederate not in HSM hybrid mode



PingFederate in HSM hybrid mode





PingFederate supports the following HSMs:

- AWS CloudHSM (stores private keys only)
- Gemalto SafeNet Luna Network HSM (stores private keys only)
- nCipher nShield Connect HSM (stores both certificates and private keys)

When all relevant keys and certificates are stored on the HSM, administrators can turn off the HSM hybrid mode. When the HSM hybrid mode is disabled, PingFederate delegates the management of the relevant keys and certificates to the HSM.



Important:

Once the HSM hybrid mode is disabled, for keys and certificates that should be stored on an HSM, PingFederate will only access those keys and certificates from the HSM, regardless of whether such keys and certificates exist on the local trust store.

Steps

1. Install and configure the HSM client and the existing PingFederate environment. See Supported hardware security modules on page 169.



Important:

When editing the <pf install>/pingfederate/bin/run.properties file, set the pf.hsm.hybrid property to true to enable the HSM hybrid mode.

Result:

Once PingFederate is integrated with your HSM, you can create (and store) new certificates on your HSM. Because the HSM hybrid mode is enabled, you can reconfigure connections or other configuration items to use the new certificates over a period of time. As long as the HSM hybrid mode is enabled, PingFederate can use certificates stored on your HSM and the local trust store.



Important:

When making changes to keys and certificates, you might need to coordinate with your partners. For more information, see Digital signing policy coordination on page 88.

2. Create a new SSL server certificate on your HSM and activate it for the administrative console and the runtime server on Security # Certificate & Key Management # SSL Server Certificates.



Note:

You can also create separate certificates on your HSM and activate one certificate for the administrative console and the other certificate for the runtime server. For configuration steps, see Manage SSL server certificates on page 601.

3. Create new digital signing certificates and decryption keys on Security # Certificate & Key Management # Signing & Decryption Keys & Certificates and reconfigure connections or configuration items to use the new certificates and keys from your HSM.



Use **Check Usage** to locate the applicable connections or configuration items.

For configuration steps, see Manage digital signing certificates and decryption keys on page 608.

4. If your connections support outbound (SOAP) back-channel authentication by client certificates, create new SSL client certificates on Security # Certificate & Key Management # SSL Client Keys & Certificates and reconfigure connections to use the new certificates from your HSM.



Tip:

Use **Check Usage** to locate the applicable connections or configuration items.

For configuration steps, see Manage SSL client keys and certificates on page 606.



Tip:

Use **Check Usage** to locate the applicable configuration items.

For configuration steps, see *Manage trusted certificate authorities* on page 600.

6. If you are transitioning to an nCipher HSM, for connections using the unanchored trust model, export the partner certificate for back-channel authentication from the local trust store, import them to your HSM, and reconfigure the connections to use the new certificates from your HSM. For information about the unanchored trust model, see "Trust models" under Digital signing policy coordination on page 88.

For configuration steps, see *Managing certificates from partners* on page 622.

Manage Partner metadata URLs

Use Security # Certificate & Key Management # Partner Metadata URLs to add, update, review, or remove SAML metadata URLs provided by your partners.

SAML metadata URLs streamline the process of establishing and maintaining SAML connections. If your partner provides SAML metadata by URL, you can use the metadata URL for the following scenarios:

- Creating a new SAML connection using the metadata URL and associating the metadata URL with the new connection
- Enabling or disabling automatic updates from the associated metadata URL
- Adding or updating the metadata URL associated with an existing SAML connection
- Updating an existing SAML connection using the metadata URL instantly



You can quickly create connections with InCommon participants, update the connections automatically or manually as the InCommon participants update their metadata, and do so securely knowing PingFederate only commits changes to your connections after validating the digital signatures of the signed metadata.

When PingFederate accesses a digitally signed metadata URL for the first time, it validates the digital signature and stores the metadata URL and its verification certificate if the signature is correct. When an existing metadata URL is accessed, PingFederate validates the digital signature using the stored certificate. If there is a digital signature error, PingFederate aborts the process and provides an error with a recommended course of action. You can bypass the signature verification process.

Adding a new metadata URL

Use the Partner Metadata URLs window's functionality to add a custom-configured metadata URL.

Steps

- 1. On the Partner Metadata URLs window, click Add New URL.
- 2. On the **URL** tab, define the metadata URL.
 - a. Configure each field.

Field	Description
Name	A name of the metadata URL.
URL	The metadata URL.

Field	Description
Validate Metadata Signature	Determines whether PingFederate should validate the digital signature of signed metadata.
	Select the check box to verify digital signatures.
	Clear the check box to skip the signature verification process.
	This check box is selected by default.

- b. Click Load Metadata.
- 3. On the **Certificate Summary** tab, review the certificate information.



This is shown and applicable only when the Validate Metadata Signature check box on the URL tab

- If the metadata is not digitally signed (unsigned), click Verify to confirm that the unsigned metadata is reachable at the time of the configuration.
- If the metadata is signed but the certificate is provided outside of the metadata, click Import to upload the verification certificate.
- 4. On the **Summary** tab, review the configuration. Click **Done** and **Save**.

Updating an existing metadata URL

Use the Partner Metadata URLs window's functionality to update and correct the configuration of existing metadata.

Steps

- 1. On the **Partner Metadata URLs** window, select the applicable metadata by its name.
- 2. On the URL tab, update the name, URL, or digital signature verification option. Click Next.
- 3. On the Certificate Summary tab, click Verify to confirm that the unsigned metadata is reachable at the time of the configuration or update the verification certificate of a signed metadata.



Note:

This is shown and applicable only when the Validate Metadata Signature check box on the URL tab is selected.

- 4. If the metadata is signed but the certificate is provided outside of the metadata, click Import to upload the verification certificate. Click Next.
- 5. On the **Summary** tab, review the configuration, then click **Done** and **Save**.

Reviewing a metadata URL usage

Use the Partner Metadata URLs window's functionality to look over a piece of metadata's information.

Steps

1. On the Partner Metadata URLs window, select Check Usage for the applicable metadata.



The Check Usage option is shown and applicable only when the metadata is used by at least one connection.

- 2. Review the information in the pop-up window.
- 3. When finished, close the pop-up window.

Removing a metadata URL

Use the Partner Metadata URLs window's functionality to delete an unwanted piece of metadata.

Steps

1. On the **Partner Metadata URLs** window, select **Delete** for the applicable metadata.



Note:

The **Delete** option is shown and applicable only when the metadata is not used by any connections.

To cancel the removal request, select **Undelete** for the certificate.

2. Click **Save** to confirm your action.

Rotating system keys

Use the System Keys window to manually rotate your system keys to optimize your environment's security.

About this task

System keys are used in cryptographic operations to generate and consume internal tokens. These tokens are leveraged in multiple use cases such as one-time links for self-service password reset and email ownership verification. Periodic rotation ensures optimal security of your environment.

Steps

- 1. Go to Security # Certificate & Key Management # System Keys.
- 2. To rotate the system keys, click Rotate.
- 3. Click Save.

Result

PingFederate generates a new Pending key. The key that was Pending becomes the Current key. The key that was Current becomes the Previous key.

System integration

You can configure PingFederate to act in accordance with your circumstances and desires.

If PingFederate acts in a service provider (SP) role, you can configure redirect validation rules to ensure valuable information, such as user attribute values, is sent only to a list of designated target resources. If PingFederate is deployed behind a reverse proxy or a load balancer, you can configure whether and how PingFederate should extract contextual information from the requests. You can manage the availability and authentication requirements for the supporting services that PingFederate offers. You can find these menu items in the **Security** tab under the **System Integration** section.

Configuring redirect validation

Ensure that a designated target exists by validating single sign-on (SSO), single logout (SLO) and selfservice user account management transactions.

About this task

You can configure several service provider (SP) adapters to pass security tokens or other user credentials from the PingFederate SP server to the target resource via HTTP query parameters, cookies, or POST transmittal. In all cases, these transport methods carry the risk that a third party (with specific knowledge of the identity provider (IdP), the SP, or both, PingFederate endpoints; and PingFederate configuration) could obtain and use valid security tokens to gain improper access to the target resource.

This potential security threat involves using a well-formed SSO or SLO link to start an SSO or SLO request for a resource at the SP site. However, the target resource designated in the link intercepts the security token by a redirection to a malicious website. This same threat also applies to self-service user account management endpoints when such requests include the TargetResource parameter.

To prevent such an attack, PingFederate provides a means of validating SSO, SLO, and self-service user account management transactions to ensure that the designated target resource exists through a list of configurable URLs. At minimum, an expected resource requires a domain name (or an IP address) and the selection of one or more applicable request types.



Note:

The following default target URLs are always allowed, and you don't need to enter them into the list manually:

- The default target URL for any IdP connections (see Configuring default target URLs on page 694)
- The default target URL for any adapter-to-adapter mappings (see Configuring a default target URL (optional) on page 374)
- The SP default URL for successful SSO (see Configuring default URLs on page 656)
- The IdP default URL for successful SLO (see Configuring a default URL and error message on page 395)

PingFederate can also validate the error resource parameter. For more information about the InErrorResource parameter, see IdP endpoints on page 1133, SP endpoints on page 1138 and System-services endpoints on page 1151.



Important:

PingFederate enables both target resource validation and error resource validation by default in new installations.

For backward compatibility, PingFederate upgrade tools do not enable these options if they aren't selected in the previous PingFederate installation. Although optional, we strongly recommend enabling validation for both target and error resources and entering all expected resources (including the HTTPS option) to prevent unauthorized access.

Steps

- 1. Go to Security # System Integration # Redirect Validation.
- 2. Configure target resource validation options.

Option	Description
SSO	When selected, PingFederate validates the requested target resource for IdP connections, adapter-to-adapter mappings, and SAML 2.0 IdP Discovery against a list of configurable resources.
	This check box is selected by default in new installations.
	Clear the check box to disable the feature.

3. Configure error resource validation.



Select the Enable InErrorResource Validation check box to validate the requested InErrorResource parameter value against a list of configurable resources.

This check box is selected by default in new installations.

Clear the check box to disable the feature.

- 4. Define a list of expected resources.
 - a. Indicate whether to mandate secure connections when this resource is requested under Require HTTPS.



Important:

This selection is recommended to ensure that the validation will always prevent message interception for this type of potential attack, under all conceivable permutations.

This check box is selected by default.

b. Enter the expected domain name or IP address of this resource under Valid Domain Name.

Enter a value without the protocol, such as example.com or 10.10.10.10.

Prefix a domain name with a wildcard followed by a period to include subdomains using one entry. For instance, *.example.com covers hr.example.com or email.example.com but not example.com, the parent domain.



Important:

While using an initial wildcard provides the convenience of allowing multiple subdomains using one entry, consider adding individual subdomains to limit the redirection to a list of known hosts.

c. Optional: Enter the exact path of this resource under Valid Path.

Start with a forward slash, without any wildcard characters in the path. If left blank, any path under the specified domain or IP address is allowed. This value is case-sensitive. For instance, /inbound/Consumer.jsp allows /inbound/Consumer.jsp but rejects /inbound/ consumer.jsp.

You can allow specific query parameters with or without a fragment by appending them to the path. For instance, /inbound/Consumer.jsp?area=West&team=IT#ref1001 matches /inbound/Consumer.jsp?area=West&team=IT#ref1001 but not /inbound/ Consumer.jsp?area=East&team=IT#ref1001.

d. Optional: Select the check box under Allow Any Query/Fragment to allow any query parameters or fragment for this resource.

Selecting this check box also means that no query parameter and fragment are allowed in the path defined under Valid Path.

This check box is not selected by default.

- e. Select one or more request types for this resource.
 - Select the check box under TargetResource for SSO if this is an expected SSO target resource for one or more IdP connections, adapter-to-adapter mappings, or SAML 2.0 IdP Discovery.
 - Select the check box under TargetResource for SLO and Other if this is an expected target resource for SLO and self-service user account management requests.
 - Select the check box under InErrorResource if this is an expected InErrorResource parameter value.

These check boxes are not selected by default.

f. Click Add.

Use the **Edit**, **Update**, and **Cancel** workflow to make or undo a change to an existing entry. Use the **Delete** and **Undelete** workflow to remove an existing entry or cancel the removal request.

g. Repeat these steps to define multiple expected resources.



Note:

The display order does not matter. A more specific match is considered a better match and an exact match is considered the best match.

5. Click Save.

Managing partner redirect validation

PingFederate enables you to validate a parameter for single logout (SLO) in order to prevent unauthorized access.

About this task

Some of the parameters used to perform redirection represent locations at a partner site—for example, the wreply parameter in WS-Federation. To protect against session token hijacking through open redirections, PingFederate provides an option to validate wreply for single logout (SLO). Once enabled, the parameter value is managed within the connection on a per-partner basis. PingFederate amalgamates the entries from all active WS-Federation connections and validates wreply against the consolidated list.



Important:

PingFederate enables wreply validation for SLO by default in new installations.

For backward compatibility, PingFederate upgrade tools do not enable this option if it was not selected in the previous PingFederate installation. Although optional, enabling wreply validation for SLO and specifying the allowed domains and paths for each WS-Federation connection can prevent unauthorized access.

Steps

1. Go to Security # Redirect Validation # Partner Redirect Validation.

2. Select the Enable wreply Validation For SLO check box to enable this feature.



This check box is selected by default in new installations. Clear the check box to disable the feature.

3. Click Save.

Configure incoming proxy settings

Use the options in the Incoming Proxy Settings window to enable PingFederate to access information to construct correct responses to incoming requests.

When PingFederate is deployed behind a reverse proxy (or a similar network traffic management solution, such as a load-balancer), the following options enable PingFederate to use information in HTTP headers added by the reverse proxy to construct correct responses. These options, configurable on **Security** # System Integration # Incoming Proxy Settings, apply globally to all incoming requests.

HTTP header for client IP addresses

The HTTP Header for Client IP Addresses field allows you to globally specify the header name (for example, X-Forwarded-For) where PingFederate should attempt to retrieve the client IP address in all HTTP requests sent to PingFederate. Defining this field helps PingFederate identify the correct client IP address when PingFederate is operating behind a reverse proxy or load balancer.



By design, the X-Forwarded-For header exposes privacy-sensitive information, such as the IP address of the client. This header is untrustworthy when no trusted reverse proxy exists between the client and server. assumes that a trusted proxy verifies any headers configured in the Incoming Proxy Settings rather than coming directly from the client. Therefore, we recommend that customers only configure headers in **Incoming Proxy Settings** that are proxied through a trusted source.

Proxies commonly append the IP address from an incoming request to the X-Forwarded-For (or similar) header. If you enter X-Forwarded-For as the value of the HTTP Header for Client IP Addresses field, PingFederate combines multiple comma-separated header values in the same order that they are received. Define which IP address you want to use in the list box:

- Leave the default of Use Last Value to use the last value in the combined list.
- Select Use First Value to use the first value in the combined list.

HTTP header for hostname

The HTTP Header for Hostname field allows you to globally specify the header name (for example, X-Forwarded-Host) where PingFederate should attempt to retrieve the hostname and port in all HTTP requests sent to PingFederate. Proxies commonly append the hostname and port from an incoming request to the X-Forwarded-Host (or similar) header. If you enter X-Forwarded-Host as the value of the HTTP Header for Hostname field, PingFederate combines multiple comma-separated header values into the same order that they are received. Define which hostname you want to use in the list box:

- Leave the default of Use Last Value to use the last value in the combined list.
- Select Use First Value to use the first value in the combined list.

Client certificate header name and chain header name

If you use mutual client certificate authentication and want to use the Apache HTTP Server with mod ssl as the incoming proxy, configure the Apache HTTP Server to pass client certificates as HTTP request headers and enter the header names on the Incoming Proxy Settings window.

The following examples shows the Apache HTTP Server configured to pass the client leaf certificate and up to four intermediate certificates as headers.

```
SSLOptions +ExportCertData
RequestHeader set LEAF CERT "%{SSL CLIENT CERT}s"
RequestHeader set CHAINO "%{SSL_CLIENT_CERT_CHAIN_0}s"
RequestHeader set CHAIN1 "%{SSL_CLIENT_CERT_CHAIN_1}s"
RequestHeader set CHAIN2 "%{SSL_CLIENT_CERT_CHAIN_2}s"
RequestHeader set CHAIN3 "%{SSL_CLIENT_CERT_CHAIN_3}s"
```

Note:

This configuration snippet is for demonstration purposes only.

To configure PingFederate to consume these HTTP request headers for the purpose of mutual client certificate authentication:

- Enter LEAF CERT as the Client Certificate Header Name.
- Enter CHAIN as the Client Certificate Chain Header Name.



Note:

Do not enter the trailing number from the chain header names.



CAUTION:

Since HTTP request headers could potentially be forged, you should only specify a Client Certificate Header Name and a Client Certificate Chain Header Name if the Apache HTTP Server is immediately in front of your PingFederate environment. The specified values must match the header names used in the Apache HTTP Server configuration, omitting the trailing number from the chain header names.

Incoming proxy terminates HTTPS connections

The **Incoming proxy terminates HTTPS connections** option allows you to globally specify that connections to the reverse proxy are made over HTTPS even when HTTP is used between the reverse proxy and PingFederate.

Configuring service authentication

Administrators with the **Admin** role can activate and configure authentication for Attribute Query, Java Management Extensions (JMX), and SSO Directory Service.

About this task

If you are using the SAML 2.0 Attribute Query profile as a service provider (SP), then the requesting applications at your site must authenticate to the PingFederate server. For more information, see Attribute Query and XASP on page 43 and the /sp/startAttributeQuery.ping on page 1142 SP application endpoint.

Authentication is required to access PingFederate runtime data via JMX (see Runtime monitoring using JMX on page 865) or to make SOAP calls to the Connection Management Service. Authentication is optional for the SSO Directory Service. For more information, see Web service interfaces and APIs on page 1086 and SSO Directory Service on page 1090.

Note:

To help ensure network security, access to all of these services is deactivated when PingFederate is first installed.

To activate and configure authentication for the Connection Management Service, grant the administrators all three administrative roles: Admin, Crypto, and User Admin. For more information, see Connection Management Service on page 1087.

Steps

- To enable a service:
 - a. On Security # System Integration # Service Authentication, select Action # Activate for your desired service.
 - b. Enter or modify) the service account ID and define or reset the Shared Secret. You and the application developer must agree to these values.



Z Tip:

Authentication is optional for the SSO Directory Service.

 To disable a service, on Security # Service Authentication, select Deactivate under "Action" for your desired service.



Note:

Although not accessible when deactivated, the Connection Management Service and the SSO Directory Service are deployed by default with PingFederate. If your organization does not plan to use one or both of these services, you can remove the following WAR file or files:

- <pf install>/pingfederate/server/deploy2/pf-mgmt-ws.war for the Connection Management Service
- f install>/pingfederate/server/deploy/pf-ws.war for the SSO Directory Service

Account lockout protection

Account lockout protection provides a level of security to the user and can operate in multiple ways based on the PingFederate environment.

Account lockout protection prevents user accounts from locking at the underlying user repository based on too many failed authentication attempts. It also adds a layer of protection against brute force and dictionary attacks because the user is locked out for a time period when the number of failed attempts exceeds the threshold. This protection is enabled in many areas of PingFederate, including the HTML Form Adapter, the Username Token Processor, the OAuth resource owner password credentials grant type, and the native authentication scheme for the administrative console and API.



Note:

The HTML Form Adapter and the Username Token Processor provide a per-instance setting for the maximum number of failed attempts such that administrators can use unique values for different instances of the adapter or the token processor.

In a PingFederate clustered environment, depending on the chosen runtime state-management architecture, the account locking-state information is shared across a replica set, multiple replica sets, or all nodes in the cluster.

com.pingidentity.common.security.AccountLockingService.xml configuration file, located in the cpf install/pingfederate/server/default/data/config-store directory.

Configuring account lockout protection

Use PingFederate's functionality to customize your account lockout protection settings.

Steps

1. Edit the com.pingidentity.common.security.AccountLockingService.xml file, located in the the f install/pingfederate/server/default/data/config-store directory.

For more information, see the inline comments and the following table.

Property	Description
MaxConsecutiveFailures	The maximum number of failed attempts before a user is locked out for a time period.
	The default value is 3.
	Note:
	The per-instance setting in the HTML Form Adapter and the Username Token Processor overrides this property.
LockoutPeriod	The amount of time in minutes that a user is locked out when the MaxConsecutiveFailures threshold is reached.
	The default value is 1 minute.

If you have a PingFederate clustered environment, edit this file on the console node.

- 2. Save the change.
- 3. Restart PingFederate.
- 4. If you have a PingFederate clustered environment, click **Replicate Configuration** in **System # Server** # **Cluster Management**.

Password spraying prevention

Use password spraying prevention to mitigate against attacks which exploit weak or compromised passwords.

Password spraying prevention adds a layer of defense against the attack pattern where bad actors try to gain access to protected resources by using the same password, typically weak or compromised, against multiple accounts from multiple locations. When enabled, PingFederate tracks the number of failed login attempts per password. When the number of failures for a particular password reaches a threshold, that password is temporarily locked out. Password spraying prevention applies to the HTML Form Adapter, the Username Token Processor, and the OAuth 2.0 resource owner password credentials grant type.

While password spraying prevention can help mitigate the risk of unauthorized access, we recommend that you also enforce a good password policy and a multifactor authentication solution, such as PingID, to protect your organization from password spraying attacks.

In a PingFederate clustered environment, depending on the chosen runtime state-management architecture, state information is shared across a replica set, multiple replica sets, or all nodes in the cluster.

Settings for password spraying prevention are stored in the

com.pingidentity.common.security.AccountLockingService.xml configuration file, located in the cpf install>/pingfederate/server/default/data/config-store directory.

Configuring password spraying prevention

Configure how password spraying prevention functions within your PingFederate environment to customize your login security experience..

Steps

1. Edit the com.pingidentity.common.security.AccountLockingService.xml file, located in the $<\!pf_install\!>\!$ /pingfederate/server/default/data/config-store directory.

For more information, see the inline comments and the following table.

Property	Description
DoPasswordLocking	Enable (true) or disable (false) password spraying prevention.
	The default value is false.
MaxPasswordAttempts	The maximum number of failed attempts before a password is locked out for a time period.
	Applicable only if password spraying prevention is enabled.
	The default value is 5.
PasswordLockoutPeriod	The amount of time in minutes that a password is locked out when the MaxPasswordAttempts threshold is reached.
	Applicable only if password spraying prevention is enabled.
	The default value is 5 minutes.

If you have a PingFederate clustered environment, edit this file on the console node.

- 2. Save the change.
- 3. Restart PingFederate.
- 4. If you have a PingFederate clustered environment, click **Replicate Configuration** on **System** # **Server** # **Cluster Management**.

Implementing a MasterKeyEncryptor using AWS KMS

During initial startup, PingFederate automatically generates a randomized master key, which by default is not encrypted. If you are running in Amazon Web Services (AWS), you can configure PingFederate to use Amazon Key Management Services (KMS) to encrypt the master key.

Before you begin

- Make sure that you have an active connection to AWS.
- Use AWS KMS to generate a key to use for the PingFederate master key encryption.
- See https://docs.aws.amazon.com/kms/latest/developerguide/overview.html for general information about how you can manage access rights to your keys using key policies or AWS Identity and Access Management (IAM).

About this task

To configure the encryption of the PingFederate master key, modify two files: hivemodule.xml and com.pingidentity.crypto.jwk.MasterKeySet.xml.

Steps

- 1. Stop PingFederate.
- Open <pf_install>/pingfederate/server/default/conf/META-INF/hivemodule.xml in a text editor.
- 3. Locate the following lines near the bottom of the file.

4. To enable master key encryption using AWS KMS, replace the lines shown in step 3 with the following lines.

- 5. Save and close the file.
- 6. Open com.pingidentity.crypto.jwk.MasterKeySet.xml in a text editor.

The contents of the file are shown here.

7. Uncomment the <con:item name="keyId"> attribute and specify the key that you generated using AWS KMS. For example, after you've made the change, the file might look like the following.

- 8. Save and close the file.
- 9. Start PingFederate.

Result

After configuring and starting PingFederate, the PingFederate master key file, pf.jwk, is encrypted.

As an administrator, you can enable certain self-service applications to let end users better manage their accounts and, by extension, lower identity management costs.

PingFederate provides various self-service applications for end users to manage their accounts. These optional capabilities lower the costs of identity management by freeing administrators from round-theclock service requests to change passwords, reset passwords, unlock accounts, and recover usernames. Designed for ease of deployment, these capabilities are integrated into the HTML Form Adapter. Administrators can easily enable some or all capabilities with a few configuration changes on a per-adapter basis. Like other user-facing windows, you can customize and localize the user-facing templates to provide the desired user experience.

PingFederate also allows users to unlock their accounts without submitting a ticket to the IT department. When enabled with SSPR, if an account is locked, a user can initiate an account unlock request at the Sign On window or the per-adapter Password Reset endpoint. Through the HTML Form Adapter, PingFederate prompts the user to prove ownership of the account using the password reset flow.

When users succeed in proving account ownership, they are allowed to retain their current passwords or to reset their passwords as needed. Furthermore, self-service account unlock is only compatible with PingDirectory and Microsoft Active Directory. If the underlying datastore is connected to Oracle Unified Directory or Oracle Directory Server, users can only unlock their account by changing their current password through the password reset flow.



Similarly, when configuring customer identity and access management use cases, administrators can enable end users to manage their local accounts, connect or disconnect one or more social connections, and change or set the password for their local accounts. For more information, see Customer IAM configuration on page 324 and Enabling profile management on page 348.

Configuring self-service password management

Create or modify an instance of the HTML Form Adapter to enable a customized self-service password management capability.

About this task

PingFederate offers self-service username password management for users to change their network password. This optional capability is integrated into the HTML Form Adapter and the LDAP Username Password Credential Validator (PCV). You can configure PingFederate to generate notification messages when users successfully change the password associated with their accounts through the HTML Form Adapter or when their passwords are about to expire.

If you are validating credentials through the PingOne for Enterprise Directory PCV, you can also enable the change password capability. Notifications for change password and password expiry are not supported at this point.

Steps

1. On Authentication # Integration # IdP Adapters, create a new HTML Form Adapter instance.

You can also reuse an existing HTML Form Adapter instance. If you do, skip to *step 1c* to configure your adapter instance to enable self-service password management.

- a. Select the PCV instance as the credential validator.
- b. Optional: Update any default values or options.
- c. Select the Allow Password Changes check box.
- d. Configure your adapter instance options. For more information, see the following table.

Option	Effects	
Change Password Notification	Select if you want PingFederate to generate a notification message for the user who has successfully changed their password through the HTML Form Adapter.	
	Note: The message is sent to the user's email address, specifically the mail attribute value returned by the LDAP Username PCV instance.	
Show Password Expiring Warning	Select if you want the Sign On window to warn the user about an approaching password expiration.	
Change Password Notification	Select to choose a notification publisher instance.	
	Note: If you have not yet configured the desired notification publisher instance, click Manage Notification Publishers.	
Show Advanced Fields	Click to review or modify default values related to the change password capability. For example, update the Change Password Template field if you want to use a custom template to render the Change Password window.	

2. Optional: Customize and localize the on-window messages and notification messages.

Result

You have created a new instance or modified an existing instance of the HTML Form Adapter with the self-service password management capability.

When a user signs on through this adapter instance, the user has the option to change the password associated with the account using the **Change Password** link.

You can also provide your users the per-adapter Change Password endpoint /ext/pwdchange/ Identify, which allows them to change their password through this HTML Form Adapter instance without submitting single sign-on (SSO) requests.

PingFederate offers self-service password reset for users to recover their accounts if they forgot their passwords.

About this task

Integrated into the HTML Form Adapter and Password Credential Validator (PCV) framework, users reset their passwords through one of the following mechanisms:

- Authentication policy
- One-time link through email
- One-time password through email
- One-time password through text message
- PingID

The self-service password reset capability relies on the HTML Form Adapter and the LDAP Username PCV to query the required attributes for the chosen reset mechanism. PingFederate supports PingDirectory, Microsoft Active Directory, Oracle Unified Directory, and Oracle Directory Server out-of-the-box. Custom PCV implementations can also be developed to offer the self-service password reset features for users stored in non-LDAP data sources. For more information, see the ResettablePasswordCredential interface in Javadoc.



The Javadoc for PingFederate is located in the <pf install>/pingfederate/sdk/doc directory.

Steps

1. Create a new LDAP datastore. For instructions, see *Configuring an LDAP connection* on page 871. You can also reuse an existing LDAP datastore connection.



Important:

- When connecting to an Active Directory (AD) LDAP server, you must secure the datastore connection using LDAPS. AD requires this level of security to allow password changes.
- When connecting to PingDirectory, Oracle Unified Directory, or Oracle Directory Server, configure proxied authorization for the service account on the directory server. See *Proxied authorization* on page 877.
- 2. Create an LDAP username password credential validator. For instructions, see Configuring the LDAP Username Password Credential Validator on page 906.

The advanced fields on the Instance Configuration tab allow you to configure self-service password reset, account unlock, and user name recovery through the HTML Form Adapter

- 3. Create a new HTML Form Adapter instance. For complete field descriptions, see *Configuring an HTML Form Adapter instance* on page 288 and *HTML Form Adapter advanced fields* on page 298.
 - a. Go to Authentication # Integration # IdP Adapters.
 - b. In the IdP Adapters window, click Create New Instance.
 - c. On the **Type** tab, enter an instance name and ID.
 - d. From the Type list, select HTML Form IdP Adapter, and click Next.
 - e. On the **IdP Adapter** tab, click **Add a new row to 'Credential Validators'** and select the LDAP Username PCF instance defined in *step 2*.
 - f. Select the Allow Password Changes check box.
 - g. Select the **Change Password Notification** check box if you want PingFederate to generate a notification message for a user who has successfully changed their password through the HTML Form Adapter.
 - The message is sent to the user's email address, specifically the mail attribute value returned by the LDAP Username PCV instance.
 - h. Select a **Password Reset Type**. See the following table for more information.

Field

Description

Password Reset Type

Select one of the following methods for self-service password reset.

Authentication Policy

Based on the policy contract selected from the **Password Reset Policy Contract** list, PingFederate finds the applicable authentication policy to handle self-service password reset requests. If the users are able to fulfill the authentication requirements as specified by the policy, PingFederate allows the users to reset their password.

Email One-Time Link

Users receive a notification with a URL to reset their password. If you have not yet configured the desired notification publisher instance, click **Manage Notification Publishers**.

Email One-Time Password

Users receive a notification with a one-time password (OTP) to reset their password.

If you have not yet configured the desired notification publisher instance, click **Manage Notification Publishers**.

PingID

Users are prompted to follow the PingID authentication flow to reset their password.

Ensure the **PingID Username Attribute** field in the selected LDAP Username PCV instance is configured; otherwise, users will not be able to reset their password.

You must also download the settings file from the PingOne for Enterprise admin portal and upload the file to the **PingID Properties** advanced field.



Important:

Do not use a method that is already part of a multi-factor authentication policy that includes a password challenge, as that would indirectly reduce that authentication policy to a single factor. For example, if users normally authenticate with a password challenge and then PingID, the self-service password reset method should not be PingID. Instead, choose the **Authentication Policy** option, select a policy contract from the **Password Reset Policy Contract** list, and configure an authentication policy for self-service password reset.

Text Message

Users receive a text message notification with an OTP to reset their password.

Ensure the **SMS Attribute** field in the selected LDAP Username PCV instance is configured; otherwise, users will not receive text message notification for password reset.

If you have not yet configured SMS provider settings in PingFederate, click **Manage SMS Provider Settings**.

None

Users cannot reset password through this HTML Form Adapter instance.

The default selection is None

Password Reset Policy If you use an authentication policy to handle SSPR requests, you must Contract select a policy contract here.

This policy contract doesn't require any extended attributes because uses this policy only to find the applicable authentication policies for password resets.



Important:

You must use a policy contract dedicated only to password reset. You can't use this policy contract for SSO anywhere else. To define a policy contract solely for password reset, click **Manage Policy Contracts**.

An authentication policy that uses this contract allows users to reset their password. The policy should use strong authentication methods to securely identify the user. To ensure that the user authenticating in the password reset flow is associated with the target account, you must map the incoming user ID into its authentication sources.

- Select the Account Unlock check box if you want to enable self-service account unlock as well.
- j. Select a notification publisher instance from the list.

If you have not yet configured the desired notification publisher instance, click **Manage Notification Publishers**.

- k. Click Show Advanced Fields to review or modify the rest of the default values related to self-service password reset. For descriptions of all advanced fields, see HTML Form Adapter advanced fields on page 298
- 4. If you selected **Authentication Policy** as the password reset type, create a new authentication policy to handle self-service password reset requests.

Generally a password reset policy must authenticate users through means other than prompting for the forgotten passwords. It should also enforce multi-factor authentication for added security. Consider the following sample use case.

You have already created an authentication policy to protect single sign-on (SSO) requests. This policy uses an HTML Form Adapter instance to validate user credentials and an instance of the PingID Adapter for multi-factor authentication. If users satisfy both authentication requirements, the policy uses a policy contract to relay user attributes to partners. To learn more about this policy configuration, see *Defining authentication policies based on group membership information* on page 248.

Like SSO, you also want to protect self-service password reset with multi-factor authentication.

Knowing your company actively manages client certificates on company devices, you have decided to use an instance of the X.509 IdP Adapter (named X.509) as the first-factor authentication source in your password reset policy. You have extended the adapter contract with a CN attribute, through which the adapter exposes the username found in the client certificate. For added security, you intend to leverage PingID as the second-factor authentication source. Per *step 3e*, you have also created a new policy contract (named SSPR APC) for the sole purpose of SSPR. At this point, you are ready to create your password reset policy.

- a. On Authentication # Policies # Policies, click Add Policy.
- b. On the **Policy** window, enter a name (and optionally a description) for the policy.
- c. Select the X.509 IdP Adapter instance.
- d. Configure each policy path out of the X.509 Adapter instance.

Fail

Select **Done**, which terminates the self-service password reset request.

For instance, if a user submits an self-service password reset request from a personal device, the request will fail because the browser on the personal device is not equipped with the company-managed client certificate issued to that user (that is only available on that user's company device).

Success

Select the same PingID Adapter instance that you have created and used in the SSO policy.

- e. Configure the incoming user ID for the PingID Adapter instance.
 - 1. Click Options to open the Incoming User ID dialog.
 - 2. Select Adapter (X.509) under Source.
 - 3. Select CN under Attribute.
 - 4. Click **Done** to close the **Incoming User ID** dialog.

For more information, see *Specifying an incoming user ID* on page 245.

f. Configure each policy path out of the PingID Adapter instance.

Fail

Select **Done**, which terminates the self-service password reset request.

Success

Select SSPR APC, which is the policy contract created solely for password reset per step



Important:

You must not reuse this policy contract for SSO elsewhere.

g. Configure the contract fulfillment for the selected policy contract.

Because the sole purpose of the selected policy contract is to route the SSPR requests through this password reset policy, the fulfillment of this contract does not matter. It is not used elsewhere. For instance, you can configure its mapping as follows.

Contract Attribute	Source	Value
subject	Text	Benign

h. Click Done and then Save.

This sample use case demonstrates the capability and flexibility that a password reset policy offers. Depending on actual use cases, you can use a different series of authentication sources to authenticate users in a secure manner. For example, if your organization manages devices using AirWatch, you can add an instance of the AirWatch Adapter as one of the authentication sources in the password reset policy. Other similar solutions include MobileIron and Microsoft Intune.

5. Optional: Customize and localize the on-screen messages and notification messages.

Result

You successfully created a new instance or modified an existing instance of the HTML Form Adapter with the SSPR and account unlock capabilities.

When a user signs on through this adapter instance, the user has the option to reset the password or unlock the account using the **Trouble Signing On** link.

Additionally, you can also provide your users the per-adapter Account Recovery endpoint /ext/ pwdreset/Identify, which allows them to reset their password or unlock their account through this HTML Form Adapter instance without submitting SSO requests.

Use PingFederate's self-service user name recovery feature to enable users to recover their lost user names through their email addresses.

About this task

PingFederate offers self-service user name recovery for users to recover their accounts through email if they forget their user names.

This optional capability is integrated into the HTML Form Adapter and the LDAP Username Password Credential Validator (PCV). PingFederate supports PingDirectory, Microsoft Active Directory, Oracle Unified Directory, and Oracle Directory Server out-of-the-box. Custom PCV implementations can also be developed to offer the same capability for users stored in non-LDAP data sources. For more information, see the RecoverableUsername interface in Javadoc.



The Javadoc for PingFederate is located in the <pf install>/pingfederate/sdk/doc directory.

Steps

1. Go to Authentication # Integration.On the IdP Adapters window, create a new HTML Form Adapter instance.

You can also reuse an existing HTML Form Adapter instance. If you do, skip to step 1c to configure your adapter instance to enable the self-service user name recovery capability.

- a. Select the HTML Form Adapter instance as the credential validator.
- b. Optional: Update any default values or options.
- c. Select the Enable Username Recovery check box.
- d. Select a notification publisher instance from the list.



Note:

If you have not yet configured the desired notification publisher instance, click Manage Notification Publishers.

e. Click Show Advanced Fields to review or modify default values related to self-service user name recovery.

Example:

Select the Require Verified Email check box if you want PingFederate to only send user name recovery email messages to users who have proven ownership of their email addresses.

2. Optional: Customize and localize the on-window messages and notification messages.

Result

You successfully created a new instance or modified an existing instance of the HTML Form Adapter with the self-service user name recovery capability.

When a user signs on through this adapter instance, the user has the option to recover the user name using the **Trouble Signing On** link.

You can also provide your users the per-adapter Account Recovery endpoint /ext/pwdreset/ Identify, which allows them to recover their user name through this HTML Form Adapter instance without submitting single sign-on (SSO) requests.

Service provider SSO configuration

You can use the PingFederate administrative console as a service provider (SP) to configure local application-integration information and to manage connections to your identity provider (IdP)-partner sites.

Only one connection is needed per partner, even if integrating more than one web application.

While you define your entity ID on the **Federation Info** tab of the **Protocol Settings** window, you can identify your organization differently through the use of virtual server IDs on a per-connection basis. For more information, see *Multiple virtual server IDs* on page 106.

Additionally, you can deploy an SP connection to bridge a service provider to one or more identity providers through one or more authentication policy contracts. For more information, see Federation hub use cases on page 99 and Federation hub and authentication policy contracts on page 103.



This topic applies to configuration settings needed for browser-based single sign-on (SSO). Although this information also applies to WS-Trust security token service (STS), if you are using PingFederate exclusively as an STS, start with WS-Trust STS configuration on page 951.

SP application integration settings

The integration of local applications with PingFederate is the essential "last-mile" configuration that allows end-users at your identity provider (IdP) partner's website to access your protected resources.

The use of application-integration kits and a robust SDK facilitates the integration of local applications with PingFederate. For more information, see SSO integration kits and adapters on page 83.

You can configure the service provider (SP) adapters that PingFederate uses to create user sessions that allow single sign-on (SSO) access to your protected resources. You can also set Default URLs to which users can be directed during SSO or single logout (SLO) and look up system endpoints that application developers at your site need to access PingFederate's SSO/SLO services.



Note:

If your PingFederate configuration enables the WS-Trust security token service (STS), you can configure plugin token generators on the Applications tab in the Token Exchange # Token Generators window. For more information, see Service provider STS configuration on page 975.

Managing SP adapters

A service provider (SP) adapter creates a local-application session for a user to provide single sign-on (SSO) access to your applications or other protected resources. You must configure at least one instance of an SP adapter to set up connections to identity provider (IdP) partners.

About this task

You can configure multiple instances of adapters, based on one or more adapters, to accommodate the varying needs of your IdP partners.

PingFederate comes bundled with OpenToken Adapter. You can deploy additional integration kits from the Ping Identity Downloads website.

Steps

1. Go to Applications # Integration # SP Adapters.

2. In the **SP Adapters** window, choose from the following options.

Option	Description
Create New Instance	Configure a new instance
<existing instance="" link=""> under Instance Name</existing>	Modify an existing instance
Check Usage	Review the usage of an existing instance
Delete or Undelete	Remove an existing instance or cancel the removal

Important:

After installing new adapter program files, you might be required to make additional configuration changes in areas such as adapter instances and connections as prompted by the administrative console.



Note:

By default, PingFederate automatically checks multi-connection errors whenever you access this window. This verifies that configured connections are not adversely affected by changes made here.

If you experience noticeable delays in accessing this window, you can disable automatic connection validation. Go to System # Server # General Settings.

Creating an SP adapter instance

About this task

The first step in creating an adapter instance is choosing an adapter type.

Steps

- 1. Go to Applications # Integration # SP Adapters.
- 2. On the SP Adapters window, click Create New Instance.
- 3. On the **Type** tab, configure the basics of this adapter instance.
 - a. Enter the instance name and ID.
 - b. From the **Type** list, select the adapter type.
 - c. Optional: From the Parent Instance list, select an existing type. If you are creating an instance that is similar to an existing instance, you might consider making it a child instance by specifying a parent. A child instance inherits the configuration of its parent unless overridden. You can specify overrides during the rest of the setup.
- 4. Click Next.

Configuring an SP adapter instance

You are presented with different configuration parameters depending on the type of adapter that you selected.

About this task

Steps

If this is a child instance, select the override check box to modify the configuration.

- If you are configuring an instance of the OpenToken SP Adapter, see Configuring an OpenToken SP Adapter instance on page 322 for configuration information.
- If you are configuring an adapter from an integration kit (including any SaaS connector), use the information in *Integrations Overview* to configure the adapter instance.

Invoking SP adapter actions

Adapters can perform configuration assistance or validation actions; for example, testing a connection to an LDAP server. Actions may also include generation of parameters that might need to be set manually in a configuration file.

Steps

On the Actions tab, complete any actions required, and click Next.

Extending an SP adapter contract

You can configure adapters with an option allowing administrators to add to the attributes required for creating usable sessions.

About this task

This feature might be needed by a legacy application that requires different authentication than other applications under the same enterprise identity-management system.



Note:

If this is a child instance, select the override check box to modify the configuration.

Steps

- 1. On the Extended Contract tab, enter the name of an attribute and click Add. Repeat this step as needed to add another attribute.
- 2. Click Next.

Identifying the target application

On the Target App Info tab, you can enter the name of the target application and the URL of the application icon.

About this task

The URL is accessible through the IdP Adapter interface, IdpAuthenticationAdapterV2, in the PingFederate Java SDK. For more information about the SDK, see SDK Developer's Guide on page 1021. Both fields are optional.



Note:

If this is a child instance, select the override check box to modify the configuration.

Steps

- 1. Optional: Enter the application name in the **Application Name** field.
- 2. Optional: Enter the URL to the application icon in the **Application Icon URL** field.
- 3. Click Next.

Reviewing an SP adapter configuration

You can review your SP adapter settings before completing the configuration.

Steps

- To keep your changes, click Save.
- To amend your configuration, click the name of the corresponding tab and then follow the configuration wizard to complete the task.
- To discard your changes, click Cancel.

Configuring target URL mapping

When you have more than one target session defined in an identity provider (IdP) connection, you must map the target URL to its target session.

About this task

When PingFederate receives a single sign-on (SSO) or single logout (SLO) request, it compares the target URL against the configured URLs until a match is found. If a match is not found, the SSO request fails.



Important:

For target URL mapping to work correctly, you must configure a target resource entry in the Security # System Integration # Redirect Validation settings. If you have not done this, follow the instructions in Configuring redirect validation on page 631.

For example, this mapping configuration might be necessary in an IdP-initiated SSO scenario that connects to multiple applications at your site. For transactions initiated at your site, this mapping is required for default situations where the target resource and the adapter instance are not specified in the SSO or SLO request. When this information is provided with the service provider (SP) request, the mapping table is ignored. For more information, see SP services on page 1139.

When bridging an identity provider to multiple service providers, for each service provider supporting the SAML IdP-initiated SSO profile, map the target URLs to the corresponding SP connection.



In this scenario, PingFederate is a federation hub for the identity provider and the service providers. For more information, see *Federation hub use cases* on page 99.

Finally, if an IdP connection is associated with one or more SP adapters, authentication policy contracts, or both, you also need to map the target URLs to their respective target session.

You manage target URL mappings on the Applications # Integration # Target URL Mapping window. The configuration process involves entering a URL and select a target session for it. See the following table for more information.

The order of mapping is significant in that the first matching mapping, from top to bottom, determines which target session receives the request. For example, if two URLs are mapped in the following order.

URL	Session Target
http:// www.example.com/ acct101/	OpenToken SP Adapter to an local training app
http://www.example.com/*	SP connection to SP SaaS

A target URL of http://www.example.com/acct101/ will be mapped to OpenToken SP Adapter to an local training app because the target matches the first mapping in the configuration.

If the order of the mappings is reversed, the same target will be mapped to SP connection to ACME SaaS because the first mapping in the new configuration, http://www.example.com/*, matches the target URL.

Steps

1. Enter a URL.

The target URLs that align with your configured target sessions. The URLs instruct the PingFederate SP server to route session-creation processing through an SP adapter instance or an SP connection.

You can use a wildcard (*) to match multiple URLs to the same target session but you can use only one wildcard (*) per URL.

If the target URL in the incoming request is not matched by the first entry in this table, subsequent entries are tried until a match is found.



Note:

PingFederate tries the next entry if a target session is not allowed based on restrictions imposed. For more information, see Restricting a target session to certain virtual server IDs on page 678.

2. Select a target type from the list.

You can only select a target type from the list when the IdP role is activated with at least one protocol for browser-based SSO.

If the IdP role is not activated or is activated without any protocol for browser-based SSO, such as SAML or WS-Federation, the **Target Type** value defaults to **SP Adapter**.

3. Select a target session from the list.

The available values depends on the chosen the **Target Type** list.

- 4. Click Add Mapping.
- 5. Repeat these steps to add multiple mappings.

Next steps

Use the up and down arrows to re-arrange the order of the mappings. Click Edit, UpdateCancel to make or undo a change to a mapping. Click **Delete** and **Undelete** to remove a mapping or cancel the removal request.

Configuring Identity Store Provisioners

PingFederate allows you to create custom identity store provisioners to bridge the inbound system for cross-domain identity management (SCIM) processing of PingFederate to your own user store. For example, you might need to create a custom identity store provisioner that works with an applicationspecific user database schema.

Using the SDK for PingFederate, you can create and test these custom identity store provisioners. For more information, see the PingFederate SDK Developer's Guide on page 1021.

To support custom attributes, you must add the schema extension and the custom attributes to the identity provider (IdP) connection. Furthermore, you need to take the expected data structure of the custom attributes into consideration when implementing the IdentityStoreProvisioner interface and its methods. In other words, your methods must be able to create, read, update, and delete/deactivate the custom attributes and their sub-attributes if the custom attributes are complex attributes to and from your user store. For more information about custom attributes, complex attributes, and other attribute types, see Defining custom SCIM attributes on page 713 and SCIM 1.1 Core Schema.



The identity store provisioner option is active only after you enable Inbound Provisioning.



Note:

By default, PingFederate automatically checks multi-connection errors whenever you access this window. This verifies that configured connections are not adversely affected by changes made here.

If you experience noticeable delays in accessing this window, you can disable automatic connection validation. Go to System # Server # General Settings.

Creating an Identity Store Provisioner instance

On the Type tab, you begin creating an instance of an identity store that PingFederate uses to bridge the inbound System for Cross-domain Identity Management (SCIM) processing to an external user store using a custom implementation.

Steps

- Go to System # Data & Credential Stores # Identity Store Provisioners.
- 2. Click Create New Instance.
- 3. On the **Type** tab, enter the a name and an ID in the **Instance Name** and **Instance ID** fields.
- 4. From the **Type** list, select a provisioner type. Available provisioner types are limited to those that are currently installed on your server.
- 5. Optional: Select a parent instance from the **Parent Instance** list. If you are creating an instance that is similar to an existing instance, you might consider making it a child instance by specifying a parent. A child instance inherits the configuration of its parent unless overridgen. You can specify overridges during the rest of the setup.
- 6. Click Next.

Defining the Identity Store Provisioner behavior

About this task

Different configuration parameters are available on the Identity Store Provisioner tab. These options are controlled by the provisioner plugin software (for more information, see topics about the Identity Store Provisioner interfaces in the SDK Developer's Guide on page 1021).

Steps

- 1. Configure the options as needed.
- Click Next.

Extending the Identity Store Provisioner contract

Identity Store Provisioners can be written with an option allowing administrators to add to the core attributes the plugin instance requires.

About this task

Both the core and extended contract attributes you define must be mapped when you configure Write Users within an inbound provisioning connection.



Z Tip:



Note:

If this is a child instance, select the override check box to modify the configuration.

Steps

1. On the **Extended Contract** tab, create and modify attributes as needed.

Option	Action
Add an attribute	Enter the attribute name in the text box and click Add .
Modify an attribute name	a. Click Edit under Action for the attribute.b. Make the change and click Update.
	Note: If you change your mind, click Cancel under Action.
Delete an attribute	Click Delete under Action for the attribute.

2. Click Next.

Extending the Identity Store Provisioner contract for groups

Identity Store Provisioners can be written with an option allowing administrators to add to the core group attributes the plugin instance requires.

About this task

Both the core and extended group attributes that you define must be mapped when you configure Write Groups within an inbound provisioning connection.



To keep your plugin flexible across multiple connections, assuming a one-to-one connection-to-identity store provider setup, you might want to hard code a set of core attributes for all connections to fulfill, and then extend attributes on as needed when a partner connection depends on additional attributes.



Note:

If this is a child instance, select the override check box to modify the configuration.

1. On the Extended Group Contract tab, create and modify attributes as needed.

Option	Action
Add an attribute	Enter the attribute name in the text box and click Add .
Modify an attribute name	 a. Click Edit under Action for the attribute. b. Make the change and click Update Note: If you change your mind, click Cancel under Action.
Delete an attribute	Click Delete under Action for the attribute.

2. Click Next.

Reviewing the Identity Store Provisioner configuration

On the **Summary** tab, review your Identity Store Provisioner configuration.

Steps

- To amend your configuration, click the corresponding tab title and then follow the steps to complete the task.
- To keep your changes, click Save.
- To discard your changes, click Cancel.

Configuring default URLs

As a service provider (SP), you can supply a default URL that the end-user might see when a single sign-on (SSO) request succeeds, that is, a session is created at your site, but the target resource is not available or not specified.

To configure default URLs, go to Applications # Integration # SP Default URLs.



Note:

You can also specify default target SSO URLs for individual identity provider (IdP) connections, which take precedence over this global setting. For more information, see Configuring default target URLs on page 694.

Similarly, you can specify to prompt a default URL indicating a successful single logout (SLO) to the enduser if no other page is designated.



Note:

The error message is displayed only when the application calling the start-SSO endpoint does not explicitly provide its own error page URL. The default entry in this field is used to localize the message. For information about how to find and change the default English message and how use the PingFederate localization feature, see Localizing messages for end users on page 817. If localization is not needed, you can also specify a message directly in this field to change the default.

Your application or your partner's application can supply these URLs at runtime (see SP services on page 1139). However, if none are provided, PingFederate uses the default values you enter on this window unless, in the case of SSO, a default is also defined for the connection.



If no default targets are specified here or at the connection level (for SSO), PingFederate provides builtin landing pages for the user. These web pages are is among the templates you can modify with your own branding or other information. For more information, Customizable user-facing pages on page 798.

Viewing SP application endpoints

Web-application developers at your site need to know the application endpoints to initiate transactions through PingFederate.

Go to Help # SP Application Endpoints to see a list of endpoints and descriptions applicable to your federation role.



Note:

These endpoints are built into PingFederate and cannot be changed.

For specific parameters required or allowed for these endpoints, see SP services on page 1139 and System-services endpoints on page 1151.

Federation settings

If your identity federation uses the SAML 2.0 X.509 attribute sharing profile (XASP), you might need to identify the identity provider (IdP) connection to which an attribute request applies.

If so, use the System # Protocol Metadata # Attribute Requester Mapping window to complete the configuration. For more information, see Attribute Query and XASP and Managing attribute requester mappings on page 657.

View endpoints that your federation partners need to know to access your services from the **Help** menu.

Managing attribute requester mappings

If you are using the SAML 2.0 X.509 attribute sharing profile (XASP), applications at your site must supply the subject distinguished name (DN) to identify a user's X.509 authentication certificate.

About this task

Optionally, an application can also supply an issuer DN, which can be used to determine the correct identity provider (IdP) attribute authority to use for a set of users associated with an IdP. For more information, see Attribute Query and XASP on page 43.



Note:

You must set the Format query parameter to a specified value for XASP. For more information, see SP services on page 1139.

You can map X.509 identifying information to connections and specify a default connection on the **System** # Protocol Metadata # Attribute Requester Mapping window.

At runtime, the issuer DN, if supplied, is evaluated against the entries under Issuer DN Pattern in hierarchical order until a match is found. If a match is found, the corresponding IdP connection is selected to issue a response to the attribute guery request. If the issuer DN matches no entry or if it is not provided, the subject DN from the request is compared against the entries under **Subject DN Pattern** in a similar manner. If the subject DN matches no entry, then the default IdP connection is used.

You can use a regular expression to match different DNs to the same connection. Only one expression can be used in any single entry. DN values must be entered in all lower-case characters.

Steps

- 1. Map one or more issuer DNs to SAML 2.0 IdP connections, as needed.
 - a. Enter an issuer DN under Issuer DN Pattern.
 - b. Select an IdP connection under IdP Connection Name.
 - c. Click Add.
 - d. Repeat these steps to add more entries.
- 2. Map one or more subject DNs to SAML 2.0 IdP connections, as needed.
 - a. Enter a subject DN under Subject DN Pattern.
 - b. Select an IdP connection under IdP Connection Name.
 - c. Click Add.
 - d. Repeat these steps to add more entries.
- 3. Select a default IdP connection from the list.

Next steps

You can click **Edit**, **Update**, and **Cancel** to make or undo a change to an entry. Click **Delete** and **Undelete** to remove an entry or cancel the removal request.

Viewing SP protocol endpoints

Your federation partners or security token service (STS) clients need to know the applicable service provider (SP) services endpoints to communicate with your PingFederate server. Configured service endpoints for SAML connections are included in metadata export files.

Go to **Help # SP Endpoints** to see a list of applicable OpenID Connect, SAML, WS-Federation, and WS-Trust STS endpoints. A pop-up window displays only those endpoints related to the federation protocols enabled on the **System # Server # Protocol Settings # Federation Info** tab. These endpoints are built into PingFederate and cannot be changed.

PingFederate provides a favorite icon for all protocol endpoints. For more information, see *Customizing the favicon for application and protocol endpoints* on page 836.

The table below describes each endpoint.

Service	URL and Description
Third Party Initiated Login (OpenID Connect 1.0)	/sp/init_login.ping
	The URL that receives and processes login requests initiated by an OpenID Provider (OP) or another party. This protocol endpoint supports HTTP GET and POST methods and the following parameters:
	 iss: the Issuer Identifier of the OP, to which PingFederate sends the authentication requests.
	This parameter is always required. • target_link_url: the destination of the request after authentication.
	This parameter is required if no default target URL is specified in the SP configuration or the IdP connection. If specified, the parameter value always overrides the default target URL. • login_hint: a hint to the OP about the end user.
	This parameter is optional.
	If your use case supports a generic login, you can add the <code>login_hint</code> parameter with a default value to the IdP connection on the <code>OpenID</code> Provider Info window. Furthermore, you may select the check box under Application Endpoint Override so that the application can optionally override the login hint value by including the <code>login_hint</code> parameter in the URL.
	Other parameters if any are not sent to the OP unless they are defined with a default value (or default values) in the IdP connection. For more information, see .
	For more information about Third Party Initiated Login flow, see the OpenID Connect specification.
Single Logout Service	/sp/SLO.saml2
(SAML 2.0)	The URL that receives and processes logout requests and responses.
Assertion Consumer	/sp/ACS.saml2
Service (SAML 2.0)	A SAML 2.0 implementation that receives and processes assertions from an IdP. The numbers reflect the index value PingFederate uses to handle each binding.
Artifact Resolution	/sp/ARS.ssaml2
Service (SAML 2.0)	The SOAP endpoint that processes artifacts returned from a federation partner to retrieve the referenced XML message on the back channel. See the note at the end of this table.
Assertion Consumer	/sp/acs.saml1
Service (SAML 1.x)	A SAML 1.x implementation URL that receives and processes assertions from an IdP.
Single Sign-on Service	/sp/prp.wsf
(WS-Federation)	The WS-Federation implementation URL that receives and processes security tokens and SLO messages.



Important:

If mutual SSL/TLS is used for authentication, you must configure a secondary PingFederate listening partner for use by partners or STS clients for the relevant endpoints such as *.ssaml* and *.wst. For more information, see .

Virtual server ID support

For SAML connections using multiple virtual server IDs, each virtual server ID has its own set of protocol endpoints. For more information, see *Multiple virtual server IDs* on page 106. You can export a connection metadata for your partner on the System # Protocol Metadata # Metadata Export window. For more information, see Exporting connection-specific SAML metadata on page 860.

For WS-Federation and SAML connections using multiple virtual server IDs, you can provide your partner the federation metadata endpoint (/pf/federation metadata.ping) with the PartnerIdpId and vsid parameters. See the table below for an example.

Partner's entity ID	Your virtual server ID	Federation metadata URL
SP	idev1	https://www.example.com/pf/sts_mex.ping? PartnerIdpId=IdP&vsid=idev1
	idev2	https://www.example.com/pf/sts_mex.ping? PartnerIdpId=IdP&vsid=idev2

In this example, the base URL and the runtime port of your PingFederate server are www.example.com and 443, respectively.

The federation metadata endpoint returns information that is specific for a given virtual server ID when the request includes the vsid parameter.

For WS-Trust STS, you can provide your partner the STS metadata endpoint (/pf/sts_mex.ping) with the PartnerIdpId and vsid parameters. The STS metadata endpoint returns information that is specific for a given virtual server ID when the STS metadata request includes the vsid parameter.

For more information about these metadata endpoints, see System-services endpoints on page 1151.

The virtual server ID concept does not apply to the /pf/sts.wst endpoint because token-to-token exchange does not involves any connections. As needed, you can pass the token-to-token endpoint to your partners as-is.

Managing IdP connections

As a service provider (SP) site, you can manage connection settings to support the exchange of federation-protocol messages, such as OpenID Connect, SAML, WS-Federation, or WS-Trust, with an identity provider (IdP), OAuth client, OpenID Provider (OP), or security token service (STS) client application at your site.

These settings include:

- User attributes that you expect to receive in an SSO token such as a SAML assertion or WS-Trust STS SAML token.
- User attributes the you expect the OP to return in an ID token or through its user information, UserInfo, endpoint on-demand.
- User attributes that may be requested using the SAML Attribute Query profile if that profile is used.
- The protocol, profiles, and bindings of the connection, including detailed security specifications such as the use of back-channel authentication, digital signatures, signature verification, and XML encryption.

To establish a connection, you and your partner must have decided this information in advance. For more information, see .

As an SP site, you respond to user requests for single sign-on (SSO) and single logout (SLO) by creating or closing user sessions, respectively, in local applications. You integrate these applications with PingFederate by configuring them with SP adapter instances. Furthermore, in preparation for configuring a new SSO connection, you need to know which adapter instance or authentication policy contract to use. For more information, see *Managing target session mappings* on page 676.

No adapter instance or authentication policy contract is required for a connection that uses only the Attribute Query profile. For more information, see *Manage the Attribute Query profile in an IdP connection* on page 698.

If you intend to pass attribute values to an adapter instance from a local datastore, you must define the datastore during this configuration. If you have not done so already, see *Datastores* on page 867.

Administrative interface

You manage connection settings in the **Authentication # Integration # IdP Connections** window, which organizes the settings into a series of primary tasks. Some primary tasks have one or more levels of sub tasks. Each primary or sub task has its own tab, where you manage one or more settings. You can move to a sibling task using the **Next** or **Previous** button. If you are on a sub task, you can also move to its parent task using the **Done** button.

When creating a new connection, you can save your progress using the **Save Draft** button. Note that not all tabs offer this option. When you reach the **Activation & Summary** tab, you must click **Save** to complete the new connection.

When editing an existing connection, you can make changes and then click **Save** to commit your changes. In order words, you are not required to step through all tabs to reach the **Activation & Summary** tab before you can save your changes.



Note:

The Save button is available on most tabs. If a tab does not show a Save button, click Next or Done until you reach to a tab where you can use the **Save** button to commit your changes.

Accessing IdP connections

In the IdP Connections window, you can create or import a connection, or edit a recently modified connection by clicking on its connection name.

About this task

The IdP Connections window displays 20 connections at a time. As needed, use the pagination controls to navigate through the rest of your connections. You can also search connections by their names or connection IDs.



A connection is included in the search results so long as its name or ID is a partial, case-insensitive match to a search term.

You can sort by connection name, partner connection ID, and default virtual server ID; narrow by protocol and status; and perform various connection-related tasks.

Steps

Go to the Authentication # Integration # IdP Connections

Choice	Action
Edit a connection	Select the connection by its name. For the setting you want to change, select the corresponding tab and follow the configuration wizard to complete the task.
Create a connection	Click Create Connection , then follow the configuration wizard to create a new connection to your identity provider (IdP) partner.
Copy a connection	Click Select action # Copy , then follow the configuration wizard to create a new connection based on an existing (source) connection.
	This is most useful if the new connection and the source connection share many common setting values.
Export a connection	Click Select Action # Export Connection , then save the XML file as prompted.
	This is useful in situations where you want to make a backup of a connection prior to making changes to it.
Import a connection	Click Import Connection , then follow the onscreen instructions to complete the task.
	If the connection already exists, you have the option to overwrite the existing connection.
	Note:

Choice	Action
	Prior to the import, you can modify the XML file to suit your needs. The XML file can also be imported to another PingFederate environment acting in the same federation role (SP) at your site. The source and the target must run the same version of PingFederate.
Export metadata for any SAML browser single sign-on (SSO) connection	Click Select Action # Export Metadata , then follow the on-window instructions to complete the task.
Update a SAML browser SSO connection	Click Select Action # Update with Metadata, then follow the on-screen instructions to complete the task. You can update a connection via a metadata XML file or a metadata URL.
	Important: The update operation might require additional configuration. Review the connection after the update operation.
toggle the status of a connection	Slide the toggle switch to enable or disable a connection.
remove a connection	Click Select Action # Delete.
Override the verbosity of runtime transaction logging for all IdP connections	Click Show Advanced Fields and the select the desired override option.
	Override Description option
	Off Select this option and let the per-connection Logging Mode configuration determine the amount of information PingFederate records in the runtime transaction log.
	This is the default selection. On Select this option, followed by one of the four logging modes, to set the verbosity of runtime transaction logging for all IdP connections. This is most useful when troubleshooting an issue that affects multiple connections.
Turn off automatic multi-connection error checking	Click Show Advanced Fields and the select the Disable Automatic Connection Validation check box.
	This check box is not selected by default.

Choice	Action
	Once selected or cleared, the state of this setting is reflected on Applications # Integration # SP Connections as well.
	For more information about this advanced setting and its impact, see <i>Configuring automatic connection validation</i> on page 790.
Keep your changes	Click Save.
Discard your changes	Click Cancel.

Resolving IdP connection errors

About this task

PingFederate automatically validates configured connections before displaying them on the IdP Connections window. This validation ensures these connections have not been adversely affected by any subsequent changes in the supporting components, such as an adapter instance or an authentication policy contract.

If errors are found, the administrative console displays a visual cue next to the applicable connections.

Steps

 To resolve the error, select the connection and follow the on-window instructions to modify the configuration, one connection at a time.

Choosing an IdP connection type

You can use the PingFederate administrative console to choose an identity provider (IdP) connection type.

About this task

You can indicate on the **Connection Type** tab whether the connection to this partner is for browser single sign-on (SSO), WS-Trust security token service (STS), OAuth, SAML, inbound provisioning, or a combination of them.



Note:

You can add STS, OAuth, and outbound provisioning support to any existing SSO connection, or vice versa, at any time. However, when OpenID Connect is the chosen protocol for browser SSO, the other types become unavailable.

Select the applicable protocol on the **Connection Type** tab when establishing a new connection.



Note:

If your partner's deployment also supports multiple protocols and you intend to communicate using more than one, you must set up a separate connection for each protocol. Each connection must use a unique partner connection ID.

• On the **Connection Type** tab, indicate the desired type of connection to your partner.

On the connection Type tab, indicate the desired	type of conficction to your partner.
Choice	Action
Configure a connection for secure browser-based SSO	Select the Browser SSO Profiles check box and a protocol from the list, if necessary.
Configure an STS connection	Select the WS-Trust STS check box and the default token type from the list.
Configure a connection that exchanges SAML	Select the OAuth Assertion Grant check box.
assertions or JSON web tokens (JWTs) for access tokens	Note:
	The OAuth Assertion Grant option is available only if at least one Access Token Manager instance has been configured on the Applications # OAuth # Access Token Management window
	For more information about these standards, see Security Assertion Markup Language (SAML) 2.0 Profile for OAuth 2.0 Client Authentication and Authorization Grants and JSON Web Token (JWT) Profile for OAuth 2.0 Client Authentication and Authorization Grants.
Configure an inbound provisioning connection	Select the Inbound Provisioning check box and choose to support provisioning of users only (User Support) or users and groups (User and Group Support). For groups, nested group membership, if any, is preserved.

• Optional: If your PingFederate license manages connections by groups, you can select a group for this connection.

This option is not displayed for unrestricted or other types of licenses.

Choosing IdP connection options

On the **Connection Options**tab, shown only for browser-based single sign-on (SSO) connections, you can enable browser-based SSO in conjunction with Just-in-Time (JIT) provisioning. Additionally, you can also choose to map user attributes for persistent grants used by the optional PingFederate OAuth authorization server.

About this task

For SAML 2.0, you can configure the Attribute Query profile with or without the browser-based SSO.

Steps

• On the Connection Options tab, make the appropriate selections for your configuration.

Choice	Action
Create a connection for browser-based SSO.	Select the Browser SSO check box.
Enable JIT provisioning, OAuth attribute mapping, or both.	Select the appropriate check box after selecting the Browser SSO check box.

Choice	Action
Create a connection to facilitate the SAML 2.0 Attribute Query profile.	Select the Attribute Query check box. For more information, see <i>Attribute Query and XASP</i> on page 43

Importing IdP metadata

You can use the PingFederate administrative console to import and update identity provider (IdP) metadata.

About this task

If you are using one of the SAML protocols without a connection template, you can expedite the setup by one of the following actions:

- Import a metadata file
- Select a metadata URL

When you select a metadata URL, PingFederate also enables the automatic update option and checks the metadata periodically. If PingFederate detects changes in the partner's signing certificates for digital signature verification, encryption key, or contact information, it updates the connection automatically. For better housekeeping, the update process removes verification certificates from the connection when the partner no longer maintains them in its metadata. In a clustered environment, PingFederate automatically replicates verification certificates and encryption key changes to all engine nodes. Offline engine nodes will also consume these changes as they restart and rejoin the cluster. If you prefer to update the connection manually, you can clear the **Enable Automatic Reloading** check box.

You can configure reload frequency at System # Protocol Metadata # Metadata Settings # Metadata **Lifetime** tab. The default reload frequency is daily.

We recommend you turn on notifications for SAML metadata update events at System # Monitoring & Notifications # Runtime Notifications.



Note:

The notification message provides a list of the applicable items if the metadata contains changes that require additional configuration.

After creating the connection, you can add, remove, or change the metadata URL associated with the connection in the Metadata URL tab. In addition, you can toggle the Enable Automatic Reloading check box for the connection.



Using a metadata URL with automatic reloading streamlines the configuration process. For example, you can quickly establish a browser SSO connection to an InCommon-participating partner. For more information, see www.incommon.org/participants.

Steps

1. Select from one of the following steps to import or update metadata.

Metadata medium	Steps
Metadata file	On the Import Metadata tab, select the File option.

Metadata medium	Ste	ps
	b.	Choose the metadata file, and then click Next .
		Note:
		If the metadata contains multiple entries, select the desired partner from the Select Entity ID list and click Next .
		Note:
		If the metadata file is digitally signed but the verification certificate is provided outside of the metadata, import the metadata verification certificate on the Import Certificate tab, and then click Next.
	C.	On the Metadata Summary tab, review the signature information to evaluate the authenticity of the metadata.
Metadata URL		On the Import Metadata tab, select the URL option. Select the metadata from the Metadata URL list.
		Tip:
		If the metadata you want is not shown in the list, click Manage Partner Metadata URLs . For more information, see <i>Manage Partner metadata URLs</i> on page 629.
	C.	Optionally, clear the Enable Automatic Reloading check box to disable automatic update.
		Note:
		A warning will display if you do not have runtime notifications enabled. To enable these notifications, go to System # Monitoring & Notifications # Runtime Notifications and select the Notification for SAML Metadata Update Events box.

Metadata medium	Steps
	d. Click Load Metadata.
	Note:
	If the metadata contains multiple entries, select the desired partner from the Select Entity ID list and click Next .
	Note:
	If there is a digital signature error, click Manage Partner Metadata URLs to resolve the issue.

2. Click Next.

Identifying the partner

When creating an identity provider (IdP) connection, you must identify your partner and provide basic information about them.

About this task

On the **General Info** tab, you provide your partner's unique federation identifier, the display name of the connection, and some other optional information, such as virtual server IDs, contact information, and logging mode for runtime transaction logging.

In addition, on this tab you can define a default error message that end users see in the event that single sign-on (SSO) fails.

Steps

1. Provide the basic information to identify your partner.

See the following table for more information.

Field	Description
Partner's Entity ID,	The published, protocol-dependent, unique identifier of your partner.
Issuer, Partner's Realm, or Connection ID	For a SAML 2.0 connection, this is your partner's SAML Entity ID. For a SAML 1.x connection, this is the audience your partner advertises. This ID might have been obtained out-of-band or through a SAML metadata file.
(Required)	For a WS-Federation connection, this is your partner's Realm.
	For an OpenID Connect connection, this is the Issuer Identifier of the OpenID Provider (OP).
	For a security token service (STS)-only connection, this ID can be any unique identifier.

Field	Description
Enable Additional Issuers	When selected, PingFederate takes into consideration additional issuers when validating ID tokens obtained through this connection.
(Applicable only to OpenID Connect connection)	Tip:
	Enable this option if you want to support multi-tenant OpenID providers, such as Microsoft Azure AD.
	This check box is not selected by default.
	Issuer information is defined on the Additional Issuers tab.
Connection Name	A plain-language identifier for the connection: for example, a company
(Required)	or department name. This name is displayed in the connection list on the administrative console.
Virtual Server IDs (Not applicable to OpenID Connect	If you want to identify your server to this connection partner using an ID other than the one you specified on the Federation Info tab, enter a virtual server ID in this field and click Add .
connections)	Enter additional virtual server IDs as needed.
Client ID and Client	The client ID and the client secret to communicate with the OP.
Secret	This client represents PingFederate and is created and managed at the
(Applicable to and required for OpenID Connect connections)	OP. For more information, see the documentation provided by the OP.
Base URL	The fully qualified hostname and port on which your partner's federation deployment runs (for example, https://www.example.com:9031). This entry is an optional convenience, allowing you to enter relative paths to specific endpoints, instead of full URLs, during the configuration process.
Company	The name of the partner company to which you are connecting.
Contact Name	The contact person at the partner company.
Contact Number	The phone number of the contact person at the partner company.
Contact Email	The email address for the contact person at the partner company.
Error Message (Applicable only to	If an error occurs on this server, the end user's browser might be redirected in a default situation to an error page hosted within PingFederate.
SAML or OpenID Connect connections)	The default entry <code>errorDetail.spSsoFailure</code> is a variable from the $\prox pinstall > \prox pingfederate / server / default / conf / language-packs / pingfederate - messages.properties file, which is part of the PingFederate localization framework. If localization is not needed, you can also specify a message directly in this field to change the default.$
Logging Mode	The level of transaction logging applicable for this connection.
	The default selection is Standard .

If the OP supports the OpenID Connect Discovery specification, the connection setup is expedited by loading the metadata from the OP. Based on the discovery specification, PingFederate makes a direct HTTP GET request to the /.well-known/openid-configuration endpoint at the OP

and populates the attribute contract and the protocol settings of the connection automatically. Manual adjustments can be made during the connection setup or at a later time.

Additionally, you can refresh the protocol settings by reloading the metadata from the OP at any time. If additional claims are supported, PingFederate adds them to the attribute contract, so that they can be mapped to the target applications later. In the event that the previously supported claims have been dropped by the OP from the metadata, they remain in the attribute contract. While the existing mapping configuration might not be adversely affected, runtime errors might occur if certain attributes are no longer available to the target applications. If runtime errors occur, review the server log and modify the mapping configuration accordingly.

2. Optional: Click Load Metadata.

This step is not applicable to an STS-only connection.

Populating extended property values for IdP connections

You can add, remove, or update extended property values.

Steps

Optional: On the **System** tab, add, remove, or update one or more values for any extended properties defined in **System** # **Server** # **Extended Properties**.

Extended property values can serve as metadata. They can also help drive authentication requirements. For more information, see *Extended properties* on page 856.

Defining additional issuers

You can define additional issuers when creating an identity provider (IdP) connection.

About this task

On the **Additional Issuers** tab, define additional issuers that PingFederate can accept when validating ID tokens obtained through this connection. This tab appears only when the **Enable Additional Issuers** check box on the **General Info** tab is selected.

Steps

- 1. Go to the Additional Issuers tab.
- 2. Optional: Select the **Accept All Issuers (Not Recommended)** check box if you want PingFederate to accept any issuers when validating ID tokens obtained through this connection.



As suggested by the property name, we do not recommend accepting any issuers.

This check box is not selected by default.

3. Optional: Define additional issuers.

Applicable only when the Accept All Issuers (Not Recommended) check box is not selected.

- a. Enter the issuer under Additional Issuer.
- b. Optional: Enter information about the issuer under **Description**.

The **Primary Issuer** field represents the issuer defined on the **General Info** tab and is always accepted.

Configure SP Browser SSO

PingFederate supports multiple configurations for browser single sign-on (SSO) with different federation standards. You can configure these options from the Browser SSO tab.

Browser single sign-on (SSO) relies on a user's web browser and HTTP requests to broker identityfederation messaging in XML or JSON web tokens (JWT) between an identity provider (IdP) and a service provider (SP). In contrast, WS-Trust security token service (STS) messaging is typically application-driven across the back channel and does not require browser mediation.



Many steps involved in setting up a federation connection are protocol-independent; that is, they are required steps for all connections, regardless of the associated standards. For more information, see Federation roles on page 28. Also, for any given connection, some configuration steps are required under the applicable protocol, while others are optional. Still others are required only based on certain selections. The administrative console determines the required and optional steps based on the protocol and dynamically presents additional requirements or options based on selections.

The following sections provide sequential information about every step you might encounter while configuring browser-based SSO, regardless of the protocol you are using for a particular connection.

SAML 2.0 configuration steps

- Selecting SAML profiles on page 672
- Configuring user-session creation on page 673
 - Choosing an identity mapping method for SP SSO on page 673
 - Defining an attribute contract on page 674
 - Managing target session mappings on page 676
- Managing protocol settings on page 684
 - Specifying SSO service URLs (SAML) on page 685
 - Defining SLO service URLs (SAML 2.0) on page 687
 - Selecting allowable SAML bindings (SAML) on page 688
 - Specifying an artifact lifetime (SAML 2.0) on page 688
 - Defining artifact resolver locations (SAML) on page 689
 - Configuring default target URLs on page 694
 - Overriding authentication context in an IdP connection on page 695
 - Configuring signature policy on page 696
 - Specifying XML encryption policy (for SAML 2.0) on page 696

SAML 1.x configuration steps

- Selecting SAML profiles on page 672
- Configuring user-session creation on page 673
 - Choosing an identity mapping method for SP SSO on page 673
 - Defining an attribute contract on page 674
 - Managing target session mappings on page 676
- Managing protocol settings on page 684
 - Specifying SSO service URLs (SAML) on page 685
 - Selecting allowable SAML bindings (SAML) on page 688
 - Defining artifact resolver locations (SAML) on page 689
 - Configuring default target URLs on page 694
 - Configuring signature policy on page 696

WS-Federation configuration steps

- Configuring user-session creation on page 673
 - Choosing an identity mapping method for SP SSO on page 673
 - Defining an attribute contract on page 674
 - Managing target session mappings on page 676
- Managing protocol settings on page 684
 - Specifying a service URL (WS-Federation) on page 686
 - Configuring default target URLs on page 694
 - Configuring signature policy on page 696

OpenID Connect configuration steps

- Configuring user-session creation on page 673
 - Choosing an identity mapping method for SP SSO on page 673
 - Defining an attribute contract on page 674
 - Managing target session mappings on page 676
- Managing protocol settings on page 684
 - Configuring OpenID Provider information on page 690
 - Configuring default target URLs on page 694
 - Overriding authentication context in an IdP connection on page 695

After configuring SSO settings, you will need to configure authentication credentials, the range of which depends on your SSO selections. For more information, see Configuring security credentials on page 731. Also, other configuration tasks might remain to be configured for new or modified connections, depending on the selected options on the Connection Options tab.

Selecting SAML profiles

A SAML profile is the message-interchange scenario that you and your federation partner agree to use. It defines the settings that support SAML usage for applications.

About this task

For SAML 2.0, PingFederate supports all identity provider (IdP) and service provider (SP)-initiated single sign-on (SSO) and single logout (SLO) profiles. For SAML 1.x, PingFederate supports both the standard IdP-initiated SSO profile and a proprietary "destination-first" SP-initiated SSO profile.



When configuring a local loopback connection, in which one PingFederate instance is both the identity provider and the service provider, disable the IdP-Initiated SLO and SP-Initiated SLO options on the Browser SSO window's SAML Profiles tab. These options determine whether SAML logout requests should be sent to the partner during the SLO flow. Those requests aren't necessary and can cause unexpected behavior when the partner connection exists locally. All local sessions for loopback connections are terminated during the SLO flow without the need to send SAML requests.

For information on typical SAML SSO and SAML 2.0 SLO profile configurations, including illustrations, see SAML 1.x profiles on page 30 and SAML 2.0 profiles on page 33.



The **SAML Profiles** tab does not apply to OpenID Connect and WS-Federation IdP connections.

Steps

Select the applicable profiles based on your partner agreement.

For SAML 2.0, you must select at least one SSO profile.

For SAML 1.x, IdP-initiated SSO is assumed and the specifications do not support SLO; the only choice on this tab is SP-initiated SSO.

Configuring user-session creation

You must create and configure user sessions when configuring service provider (SP) browser single signon (SSO).

About this task

As an SP, you must specify how PingFederate uses information sent from the IdP in SSO tokens to create user sessions for enabling access to protected resources at your site.

If you are a federation hub, bridging an identity provider to one or more service providers, you can associate one or more authentication policy contracts to the IdP connection. For more information, see .

The configuration involves choosing an identity-mapping method, establishing an attribute contract as needed, and optionally mapping one or more SP adapter instances, authentication policy contracts, or both.

Steps

On the User-Session Creation tab, click Configure User-Session Creation.

Choosing an identity mapping method for SP SSO

When configuring service provider (SP) single sign-on (SSO), PingFederate offers two methods of identity mapping you can choose from: account mapping or account linking.

About this task

PingFederate allows an SP to use either account linking or account mapping to associate remote users with local accounts for SSO between business partners. For more information, see . On the **Identity** Mapping tab, you choose which method to use in this IdP connection. You and your partner should decide in advance which option to use. For more information, see .

If your site is using account linking, then establishing an attribute contract is not required. Depending on your partner agreement, you can choose to supplement the account link with an attribute contract. In this configuration the account link is used to determine the user's identity, while the additional attributes might be used for authorization decisions, customized web pages, and so on, at the your site. For more information, see .



Important:

If you have previously set up a configuration to use an attribute contract and want to change the configuration to use account linking without additional attributes, then the existing attribute contract will be discarded.

Account linking can be used with either a clear, standard name identifier or an opaque pseudonym.

Steps

- 1. Choose which identity mapping method to use in this IdP connection. Choose from:
 - If you want to dynamically associate remote users with local accounts using a known attribute to identify a user, such as a username or email address, select Account Mapping

Account mapping uses the user identifier, SAML SUBJECT in a SAML assertion or sub in an ID token, and associated user attributes to create an association between a remote user and a local account.



If you are using PingFederate's JIT provisioning, choose Account Mapping. For more information, see Configuring just-in-time provisioning on page 699.

If you want to create a long-term association between a remote user and a local account, select **Account Linking**

To set up an attribute contract to use in conjunction with account linking, select the ... includes attributes in addition to the unique name identifier check box.



PingFederate uses a default, HSQLDB database to handle account linking. You can use your own database instead, as needed. For more information, see Account-linking datastores on page 901.



CAUTION:

Use the built-in HSQLDB only for trial or training environments. For testing and production environments, always use a secured external storage solution for proper functioning in a clustered environment.

Testing involving HSQLDB is not a valid test. In both testing and production, it might cause various problems due to its limitations and HSQLDB involved cases are not supported by Pingldentity.

- 2. If you have selected only the SP-initiated SSO profile and you intend to enforce additional authentication requirements by placing this IdP connection in an SP authentication policy, select No Mapping.
- 3. Additionally, select **No Mapping** if you are deploying an IdP connection solely for OAuth attribute mapping without the use of an authentication policy contract. For more information, see Configuring IdP connection grant mapping on page 533.

Defining an attribute contract

An attribute contract is the set of user attributes that you and your partner have agreed will be sent in single sign-on (SSO) tokens for this connection.

About this task

You can extend the attribute contract with additional attributes. Optionally, you can configure PingFederate to mask individual extended attributes in its logs. For more information, see Attribute contracts on page 93 and Attribute masking on page 95.



Z Tip:

If you are creating or updating a SAML or an OpenID Connect identity provider (IdP) connection, consider using the partner's metadata to do so. If the metadata contains the required information, PingFederate automatically populates the attribute contract for you.

Steps

1. On the Attribute Contract tab, enter the attribute name in the text box.

Attribute names are case-sensitive and must correspond to the attribute names expected by your partner.



Tip:

If you are configuring a SAML connection to an InCommon participant, the assertion might contain attributes such as urn:oid:0.9.2342.19200300.100.1.3 and urn:oid:2.5.4.42, which are standard names under various specifications, such as RFC4524 and RFC4519. For more information, see www.incommon.org/participants. The following table describes a subset of the object IDs (OIDs) referenced by the most common attributes used by InCommon participants.

OID value	Description
0.9.2342.19200300.100.1.3	mail
1.3.6.1.4.1.5923.1.1.1.1	eduPersonAffiliation
1.3.6.1.4.1.5923.1.1.1.6	eduPersonPrincipalName
1.3.6.1.4.1.5923.1.1.1.7	eduPersonEntitlement
1.3.6.1.4.1.5923.1.1.1.9	eduPersonScopedAffiliation
1.3.6.1.4.1.5923.1.1.1.10	eduPersonTargetedID
2.5.4.3	cn
2.5.4.4	sn
2.5.4.10	0
2.5.4.42	givenName
2.16.840.1.113730.3.1.241	displayName

For other attributes, see the metadata from your partner. The FriendlyName values, if available, should provide additional information about the attributes. Alternatively, third-party resources, such as www.ldap.com/ldap-oid-reference and www.oid-info.com, might help as well.

- 2. Optional: Select the check box under Mask Values in Log.
- 3. Click Add.
- 4. Repeat until all desired attributes are defined.

Next steps

Click Edit, Update, and Cancel to make or undo a change to an item. Click Delete and Undelete to remove an item or cancel the removal request.

Managing target session mappings

You can map a service provider (SP) adapter instance to an identity provider (IdP) connection and complete its mapping configuration through a series of sub tasks.

About this task

When PingFederate receives an SSO token, the corresponding SP adapter is triggered to fulfill its adapter contract based on the connection settings for the purpose of completing the "last-mile" integration with your application. As needed, you can map multiple SP adapter instances to an IdP connection, the same SP adapter instance to multiple IdP connections, or a combination of them.

Alternatively, if you use authentication policies to route users through a series of authentication sources and end each successful policy path with an authentication policy contract (APC), you can skip the mapping of an APC to an IdP connection and configure an APC-to-SP adapter instance mapping configuration.



To learn more about authentication policies, see Authentication policies on page 220.

Furthermore, you can map one or more APCs to an IdP connection to bridge an identity provider to one or more service providers. In this scenario, PingFederate is a federation hub for both sides. PingFederate uses APCs to associate this IdP connection with the applicable SP connections to the service providers; each APC has its own set of attributes to which you can map values from the SSO tokens.



Tip:

To learn more about federation hub, see Federation hub use cases on page 99.

On the Target Session Mapping tab, if presented, you must associate at least one target session, an SP adapter instance or an authentication policy contract, with an IdP connection. If you have multiple integration requirements, for example, if you are using more than one IdM system or an application not covered by a centralized system, multiple SP adapter instances. If you are bridging an identity provider to multiple service providers, map multiple authentication policy contracts.

The Target Session Mapping configuration does not apply when the No Mapping option is selected on the **Identity Mapping** tab.

Steps

On the Target Session Mapping tab:

Choice	Action
Map an SP adapter instance	Click Map New Adapter Instance.
Map an APC	Click Map New Authentication Policy.
Edit the mapping configuration of an SP adapter instance or APC	Open it by clicking on its name, select the setting that you want to reconfigure, and complete the change.
Remove an SP adapter or APC or cancel the removal request	Click Delete followed by Save or Undelete .
If you are creating a new connection and you are finished with mapping the required target sessions	Click Done .

Choice	Action
If you are editing an existing configuration and want to keep your changes	Click Save.

Result

When target sessions are restricted to certain virtual server IDs, the allowed IDs are displayed under Virtual Server IDs.



If you configure multiple target sessions for a connection, PingFederate selects the applicable adapter instance or authentication policy contract at runtime based on the target resource information in the requests and your configuration For more information, see Configuring target URL mapping on page 652.

Selecting a target session

The first step of the mapping configuration is to map an adapter instance or an authentication policy contract to your connection.

Steps

- 1. Select an adapter instance from the **Adpater Instance** list.
 - If you do not see the desired adapter instance, click Manage Adapter Instances to create a new instance of any deployed adapter.
- 2. If you want to customize adapter settings for this connection alone, select the Override Instance Settings check box.

Result:

When selected, the administrative console adds a new set of sub tasks; the Override Instance tab and its sub tasks.



Alternatively, you can create child adapter instances of a base adapter instance with overrides so that such customized settings can be applied to several connections. For more information, see Hierarchical plugin configurations on page 90.

- 3. Optional: To map an authentication policy contract (APC), select an adapter instance from the list.
- 4. Optional: If you do not see the desired APC, click Manage Authentication Policy Contracts to create a new policy contract.

Result

If you are editing a currently mapped adapter instance, you can select the **Override Instance Settings** check box. Clearing it removes all previously overridden settings for this connection. Selecting it provides you the opportunity to customize adapter settings specifically for this connection.

If you are editing a currently mapped APC, no changes can be made on this tab.

Overriding an SP adapter instance

When configuring service provider (SP) browser single sign-on (SSO), you can override the adapter instance settings.

About this task

On the **Override Instance** tab, you start a series of sub tasks to override adapter settings specifically for this connection.



Note:

Any changes to the base adapter instance are propagated to a connection provided the same changes are not overridden for the connection.

Steps

Click Override Instance Settings.

On each of the setting tabs, select the **Override** check box, make your changes, then click **Next**. When you are finished, click **Done** to continue with the rest of the mapping configuration.



Note:

The override setting tabs are functionally identical to those used for creating a new adapter instance. For more information, see *Managing SP adapters* on page 649.

Next steps

If you are editing a currently mapped adapter instance, click Override Instance Settings to reconfigure any overridden settings for this connection. You can also remove all overridden settings on a per-tab basis by clearing the **Override** check box near the top of the tab.

Restricting a target session to certain virtual server IDs

You can enforce integration requirements by restricting a target session to certain virtual server IDs when you multiplex one connection for multiple environments.

About this task

For more information, see Multiple virtual server IDs on page 106. On the Virtual Server IDs tab, follow these steps to restrict a target session to a certain virtual server ID. By default, no restriction is imposed.

Steps

- 1. Select the **Restrict Virtual Server IDs** check box.
- 2. Select the virtual server IDs that you want to allow for this target session.

Result

If you are editing a currently mapped adapter instance or authentication policy contract (APC), you can toggle the Restrict Virtual Server IDs setting. You can also change the allowed virtual server IDs.

Choosing an attribute mapping method

You can select if and how PingFederate should query a local datastore to help fulfill the attribute contract in conjunction with attribute values from the single sign-on (SSO) token.

Before you begin

To determine whether you need to look up additional values, compare the attribute contract against the adapter contract or the authentication policy contract. If the attribute contract does not contain the required information, determine whether a local datastore can supply it.

Alternatively, you can configure datastore queries as part of the fulfillment configuration for the applicable APC if you use authentication policies to route users through a series of authentication sources and end each successful policy path with an APC.

About this task

You make selections on the Adapter Data Store tab for service provider (SP) adapter mapping or the Attribute Retrieval tab for authentication policy contract (APC) mapping.



Tip:

To learn more about authentication policies, see Authentication policies on page 220.

Steps

- If the attribute contract contains all the attributes that your application requires, click Use only the attributes available in the SSO assertion.
- To set up a datastore query, click Use the SSO assertion to look up additional information, and then follow a series of sub tasks to complete the configuration. See Choosing a datastore on page 363 for step-by-step instructions.

Result

If you are editing a currently mapped adapter instance or APC, you can change the mapping method, which might require additional configuration changes in subsequent tasks.

Configuring target session fulfillment

Map values to the attributes defined for the contract. These are the values that the target application requires to create a local session for the user.

Before you begin

If you are bridging an identity provider (IdP) to one or more service providers, the values mapped to the authentication policy contracts are used by the associated service provider (SP) connections to create assertions for the service providers. For more information, see *Federation hub use cases* on page 99.

At runtime, a single sign-on (SSO) operation fails if PingFederate cannot fulfill the required attribute.

Steps

- 1. On the Adapter Contract Fulfillment tab, for each attribute, select a source from the Source list and then choose or enter a value. You must map all attributes.
 - AccountLink

When selected, the Value list populates with Local User ID. Normally, you would map Local User ID to an adapter attribute that represents the user identifier at the target. This source is not applicable to authentication policy contracts. This source appears only if you have elected to use account linking for a target session on the **Identity Mapping** tab.

Assertion or Provider Claims

When selected, the Value list populates with attributes from the SSO token. Select the desired attribute from the list.

For example, to map the value of SAML SUBJECT from a SAML assertion as the value of the subject user identifier on the contract, select Assertion from the Source list and SAML SUBJECT from the Value list.

Context

When selected, the Value list populates with the available context of the transaction. Select the desired context from the list.



Note:

As the HTTP Request context value is retrieved as a Java object rather than text, use OGNL expressions to evaluate and return values.



Note:

If you are configuring an OAuth Attribute Mapping configuration and have added PERSISTENT GRANT LIFETIME as an extended attribute in the Authorization Server Settings window, you can set the lifetime of persistent grants based on the outcome of attribute mapping expressions or the per-client Persistent Grants Max Lifetime setting.

- To set lifetime based on the per-client Persistent Grants Max Lifetime setting, select Context from the Source list and Default Persistent Grant Lifetime from the Value list.
- To set lifetime based on the outcome of attribute mapping expressions, select Expression as the source and enter an OGNL expression in the Value field.

If the expression returns a positive integer, the value represents the lifetime of the persistent grant in minutes.

If the expression returns the integer 0, PingFederate does not store the grant and does not issue a refresh token.

If the expression returns any other value, PingFederate sets the lifetime of the persistent grant based on the per-client Persistent Grants Max Lifetime setting.

To set a static lifetime, select **Text** from the **Source** list and enter a static value in the **Value** field.

This is suitable for testing purposes, or cases where the persistent grant lifetime must always be set to a specific value.

LDAP, JDBC, or Other

When selected, the **Value** list populates with attributes that you have selected from the datastore. Select the desired attribute from the list.

Expression

When enabled, this option provides more complex mapping capabilities, such as transforming incoming values into different formats. Select Expression from the Source list, click Edit under Actions, and compose your OGNL expressions. All variables available for text entries are also available for expressions. For more information, see Text.

Expressions are not enabled by default. For more information about enabling and editing OGNL expressions, see .

No Mapping

Select this option to ignore the **Value** field.

Text

When selected, the text you enter is used at runtime. You can mix text with references to any of the values from the SSO token, using the \${attribute} syntax.

You can also enter values from your datastore, when applicable, using this syntax:

\${ds.attribute}

where attribute is any attribute that you have selected from the datastore.



Two other text variables are available. **SAML SUBJECT** is the initiating user or other entity. TargetResource is a reference to the protected application or other resource for which the user requested SSO access. The \${TargetResource} text variable is available only if specified as a query parameter for the relevant endpoint, either as TargetResource for SAML 2.0 or TARGET for SAML 1.x.

You might hard-code a text value for a variety of reasons. For example, if your web application provides a consumer service, you might want to supply a particular promotion code for the partner.



Note:

If you are editing a currently mapped adapter instance or authentication policy contract (APC), you can update the mapping configuration, which might require additional configuration changes in subsequent tasks.

2. Click **Next** to continue configuration.

Defining issuance criteria for SP Browser SSO

You can define issuance criteria when configuring service provider (SP) browser single sign-on (SSO) for PingFederate.

Before you begin

Begin this optional configuration by adding a criterion. Choose the source that contains the attribute to be verified. Some sources, such as Mapped Attributes, are common to almost all use cases. Other sources, such as JDBC, depend on the type of configuration. Irrelevant sources are automatically hidden. After you select a source, choose the attribute to be verified. Depending on the selected source, the available attributes or properties vary. Finally, specify the comparison method and the desired, compared-to, value.

About this task

On the Issuance Criteria tab, define the criteria that must be satisfied in order for PingFederate to process a request further. This token authorization feature provides the capability to conditionally approve or reject requests based on individual attributes.

If you define multiple criteria, all criteria must be satisfied for PingFederate to move a request to the next phase. A criterion is satisfied when the runtime value of the selected attribute matches or does not match the specified value depending on the chosen comparison method. The multi-value contains and multi-value does not contain comparison methods are intended for attributes that might contain multiple values. Such

criterion is considered satisfied if one of the multiple values matches or does not match the specified value. Values are compared verbatim. If you require complex evaluations, including conditional criteria or partial matching, define them using attribute mapping expressions.



Important:

When you multiplex one connection for multiple environments, consider using attribute mapping expressions to verify the virtual server ID in conjunction with other conditions, such as group membership information, to protect against unauthorized access. For more information, see Multiple virtual server IDs on page 106 and Issuance criteria and multiple virtual server IDs on page 219.



All criteria defined must be satisfied or evaluated as true for a request to move forward. As soon as one criterion fails, PingFederate rejects the request and returns an error message.

Steps

1. From the **Source** list, select the attribute's source.

Depending on the selection, the Attribute Name list populates with associated attributes. See the following table for more information.

Source	Description
AccountLink	Select to evaluate the Local User ID value of the user.
	Displayed only if Account Linking is the selected identity mapping method. For more information, see .
Assertion or Provider Claims	Select to evaluate attributes from the IdP connection.
Context	Select to evaluate properties returned from the context of the transaction at runtime.
	Note: The HTTP Request context value is retrieved as a Java object rather than text. For this reason, attribute mapping expressions are more appropriate to evaluate and return values.
JDBC, LDAP, or other types of datastore (if configured)	Select to evaluate attributes returned from a data source.
Mapped Attributes	Select to evaluate the mapped attributes.

2. From the Attribute Name list, select the attribute to be evaluated.

3. From the **Condition** list, select the comparison method.

Available methods:

- equal to
- equal to (case insensitive)
- equal to DN
- not equal to
- not equal to (case insensitive)
- not equal to DN
- multi-value contains
- multi-value contains (case insensitive)
- multi-value contains DN
- multi-value does not contain
- multi-value does not contain (case insensitive)
- multi-value does not contain DN



Note:

The first six conditions are intended for single-value attributes. Use one of the multi-value ... conditions for PingFederate to validate whether one of the attribute values matches the specified value. When an attribute has multiple values, using a single-value condition causes the criteria to fail.

4. In the **Value** field, enter the comparison value.



Note:

Values are compared verbatim. If you require complex evaluations, including conditional criteria or partial matching, define them using attribute mapping expressions. For more information, see Attribute mapping expressions on page 214.

5. In the **Error Result** field, enter a custom error message.

Error results are handled in one of the two ways:

Redirect

When an InErrorResource URL is provided, the value of the Error Result field is used by the query parameter ErrorDetail in the redirect URL.

Template

When an InErrorResource URL is not provided, the value of the Error Result field is used by the variable \$errorDetail in the sp.sso.error.page.template.html template file.

Using an error code in the Error Result field allows the error template or an application to process the code in a variety of ways, for example, display an error message or e-mail an administrator.

To use localized descriptions, enter a unique alias in the Error Result field, such as someIssuanceCriterionFailed. Insert the same alias with the desired localized text in the applicable language resource files, located in the <pf install>/pingfederate/server/ default/conf/language-packs directory.

If not defined, PingFederate returns ACCESS DENIED when the criterion fails at runtime.

- 6. Click Add.
- 7. Optional: Repeat to add more criteria.

- 8. If you require complex evaluations, including conditional criteria or partial matching, define them using attribute mapping expressions. For more information, see Attribute mapping expressions on page 214.
 - a. Click Show Advanced Criteria.
 - b. In the **Expression** field, enter the required expressions.
 - c. Optional: In the **Error Result** field, enter an error code or message.



Note:

If the expressions resolve to a string value instead of true or false, the returned value overrides the Error Result field value.

- d. Click Add.
- e. Optional: Click **Test**, enter values in the applicable fields, and verify the results.
- f. Optional: Repeat to add multiple criteria using attribute mapping expressions.

Reviewing the target session mapping

You can review your target session mapping configuration and settings before completing.

Steps

- To amend your configuration, click the corresponding tab title, then follow the configuration wizard to complete the task.
- To keep your changes, click **Done** and continue with the rest of the configuration.



Tip:

When editing an existing configuration, you can also click **Save** as soon as the administrative console offers the opportunity to do so.

To discard your changes, click Cancel.

Reviewing the session creation summary

You can review your session creation settings before completing configuration.

Steps

- To amend your configuration, click the corresponding tab title, then follow the configuration wizard to complete the task.
- To keep your changes, click **Done** and continue with the rest of the configuration.



When editing an existing configuration, you can also click **Save** as soon as the administrative console offers the opportunity to do so.

To discard your changes, click Cancel.

Managing protocol settings

The **Protocol Settings** tab provides the launching point for configuring partner endpoints, message customizations, and other protocol-specific settings for browser-based single sign-on (SSO) connections.

About this task

SAML 2.0

Outbound SSO bindings (POST, redirect, artifact) and the corresponding SSO service URLs

- Outbound single logout (SLO) bindings (POST, redirect, artifact, SOAP) and the corresponding protocol endpoints
- Inbound bindings (POST, redirect, artifact, SOAP)
- Artifact lifetime
- Artifact resolution location
- Default target URL
- Authentication context mappings
- Signature policy
- Encryption policy

SAML 1.x

- Outbound SSO service URL, also known as the Intersite Transfer Service, if the service provider (SP)-Initiated SSO profile is enabled
- Inbound bindings (POST, artifact)
- Artifact resolution location
- Default target URL
- Signature policy

WS-Federation

- Protocol endpoint
- Default target URL
- Signature policy

OpenID Connect

- The scopes PingFederate sends to the OpenID provider (OP) in its authorization and token requests
- The OpenID Connect login type and authentication scheme used by PingFederate when communicating with the OP
- The authorization endpoint, the token endpoint, the user information, UserInfo, endpoint, and the JWKS URL
- Default target URL
- Authentication context mappings

Steps

To continue, click Configure Protocol Settings.

Specifying SSO service URLs (SAML)

The single sign-on (SSO) service endpoint is where PingFederate sends requests when SSO is initiated at your site according to partner requirements. It applies to all SAML versions when the service provider (SP)-initiated SSO profile is enabled.

About this task

For SAML 2.0 connections, associate bindings to the endpoints where your identity provider (IdP) partner wants PingFederate to send authentication requests when SSO is initiated at your site.

For SAML 1.x, only one endpoint is allowed, and the binding selection is not required.

Some federation use cases might require additional customizations in the authentication requests sent from the PingFederate SP server to the IdP, such as including the optional **Extensions** element in the authentication requests. You can use OGNL expressions to fulfill these use cases.

- 1. Enter an SSO service endpoint.
 - Enter the SSO service endpoint in the Endpoint URL field.

You can enter a relative path, starting with a forward slash, if you have provided a base URL on the **General Info** tab.

For SAML 1.x connections, this is the only configurable item on the SSO Service URL tab.

The remaining steps on the **SSO Service URLs** tab only apply to SAML 2.0 connections.

- b. Select a SAML binding from the list; for example, **POST**.
- c. Click Add.
- d. Optional: Repeat to add additional SSO service endpoints.
- Optional: Customize messages using OGNL expressions.

Expressions are not enabled by default. For more information about enabling and editing OGNL expressions, see Attribute mapping expressions on page 214. Additionally, message customization does not apply to SAML 1.x connection.

- a. Click Show Advanced Customizations.
- Select a message type from the list.
- c. Enter an OGNL expression to fulfill your use case.



Note:

For more information about Message Type, available variables, and sample OGNL expressions, see Customizing assertions and authentication requests on page 355.

- d. Click Add.
- e. Optional: Repeat to add another message customization.

Specifying a service URL (WS-Federation)

The service endpoint URL is a location to where PingFederate sends request for security token (RST) and single log-out (SLO) messages.

About this task

To protect against session token hijacking, PingFederate provides an option to validate wreply for SLO. When this option is enabled, you can specify additional allowed domains and paths on this tab. PingFederate validates the locations against a consolidated list of allowed domains and paths from all active WS-Federation connections before redirecting the end users to their destinations.



The settings to enter additional allowed domains and paths appear only if the option to validate wreply for SLO is enabled. For more information, see *Managing partner redirect validation* on page 634.

Steps

1. Enter the WS-Federation protocol endpoint at the identity provider (IdP) site in the Endpoint URL field.

You can enter a relative path, starting with a forward slash, if you have provided a base URL on the General Info tab.

- 2. Optional: Specify additional allowed domains and paths.
 - Indicate whether to mandate secure connections when this resource is requested under Require HTTPS.



Important:

This selection is recommended to ensure that the validation will always prevent message interception for this type of potential attack, under all conceivable permutations.

This check box is selected by default.

b. Enter the expected domain name or IP address of this resource under Valid Domain Name.

Enter a value without the protocol, such as example.com or 10.10.10.10.

Prefix a domain name with a wildcard followed by a period to include subdomains using one entry. For instance, *.example.com covers hr.example.com or email.example.com but not example.com, the parent domain.



Important:

While using an initial wildcard provides the convenience of allowing multiple subdomains using one entry, consider adding individual subdomains to limit the redirection to a list of known hosts.

c. Optional: Enter the exact path of this resource under Valid Path.

Start with a forward slash, without any wildcard characters in the path. If left blank, any path under the specified domain or IP address is allowed. This value is case-sensitive. For instance, /inbound/Consumer.jsp allows /inbound/Consumer.jsp but rejects /inbound/consumer.jsp.

You can allow specific query parameters with or without a fragment by appending them to the path. For instance, /inbound/Consumer.jsp?area=West&team=IT#ref1001 matches /inbound/Consumer.jsp?area=West&team=IT#ref1001 but not /inbound/Consumer.jsp?area=East&team=IT#ref1001.

d. Optional: Select the check box under **Allow Any Query/Fragment** to allow any query parameters or fragment for this resource.

Selecting this check box also means that no query parameter and fragment are allowed in the path defined under **Valid Path**.

This check box is not selected by default.

Defining SLO service URLs (SAML 2.0)

On the **SLO Service URLs** tab, associate bindings to the endpoints where your identity provider (IdP) receives logout requests when single logout (SLO) is initiated at your site and where PingFederate sends SLO responses when it receives SLO requests from the IdP.

About this task

This process only applies to SAML 2.0 connections when either SLO profile is selected on the **SAML Profiles** tab.

Steps

- 1. Go to Applications # Integrations # SP Connections.
- 2. Click on any SAML 2.0 connection, and then click the **Browser SSO** tab.
- 3. Click Configure Browser SSO, and then click the SAML Profiles tab.
- 4. Select a SAML binding from the list; for example, **POST**.

5. Enter the SLO endpoint URL in the **Endpoint URL** field.

You can enter a relative path, starting with a forward slash, if you have provided a base URL on the General Info tab.

6. Optional: Enter a URL in the **Response URL** field.

When specified, it is the location where SLO logout response messages are sent based on your partner agreement. When omitted, PingFederate sends logout responses to the SLO endpoint URL.

You can enter a relative path, starting with a forward slash, if you have provided a base URL on the General Info tab.

- 7. Click Add.
- 8. Optional: Repeat to add additional SLO endpoints.

Result

If you are editing an existing connection, you can reconfigure the SLO endpoints, which might require additional configuration changes in subsequent tasks.

Selecting allowable SAML bindings (SAML)

On the Allowable SAML Bindings tab, select the one or more bindings that your identity provider (IdP) partner can use to send SAML assertions or SAML 2.0 single logout (SLO) messages.

About this task

This configuration applies to all SAML connections.

Steps

Select the check boxes for only the applicable SAML bindings based on your partner agreement. Choose from:

- Artifact
- Post
- Redirect
- SOAP



Note:

If you have specified a single sign-on (SSO) or SLO endpoint using the artifact outbound binding, you must include SOAP as one of the allowable inbound binding.

Result

If you are editing an existing connection, you can reconfigure the allowable bindings, which might require additional configuration changes in subsequent tasks.

Specifying an artifact lifetime (SAML 2.0)

When PingFederate sends an artifact to your identity provider's (IdP's) single sign-on (SSO) or single logout (SLO) service endpoint, an element in the message indicates how long it should be considered valid.

About this task

On the **Artifact Lifetime** tab, specify the expiry information in seconds.

You can change the default value to meet your requirements. You should also consider synchronizing your serve clock with your partner's SAML gateway server. If clocks are not synchronized, you might need to set the artifact lifetime to a higher value to prevent latency issues.

This step applies only to SAML 2.0 connections.

Steps

Optional: Override the default value of the Artifact Lifetime field.

The default value is 60 seconds.

Result

You can update the artifact lifetime if you are editing an existing connection.

Defining artifact resolver locations (SAML)

When you enable the artifact binding as one of the allowable bindings on the Allowable SAML Bindings tab, you must provide an artifact resolution service (ARS) endpoint.

About this task

The ARS endpoint is the location where PingFederate sends back-channel requests to resolve artifacts received from the identity provider (IdP).

SAML 2.0 connections allow multiple ARS endpoints. For SAML 1.x connections, you can only enter one ARS endpoint.

Steps

- 1. Enter an ARS endpoint.
 - a. Enter the ARS endpoint URL.

You can enter a relative path, starting with a forward slash, if you provide a base URL on the General Info tab.

Result:

If you are configuring a SAML 1.x connection, you can only enter one ARS endpoint on the Artifact Resolver Location tab.

b. Optional: Enter an integer in the **Index** field for this ARS endpoint.

This is applicable only to SAML 2.0 connections.

The administrative console automatically assigns an index value for each ARS endpoint, starting from 0. If you want to define your own index values, you must make sure the index values are unique.

- c. Click Add.
- d. Optional: Repeat to add additional ARS endpoints.

This is applicable only to SAML 2.0 connections.



Note:

When specifying multiple ARS endpoints, each endpoint must share the same transport protocol. That is, if one endpoint uses HTTPS, then all must use HTTPS. Similarly, if one endpoint uses HTTP, then all must use HTTP.

2. Optional: Enter your partner's source ID.

The source ID is usually a generated value based on a federation partner's connection ID; the PingFederate service provider (SP) server will correctly generate the source ID. If that is the case for this partner, then leave this field blank. If your partner uses a Source ID that is not based on the Issuer ID, then enter the Source ID supplied by your IdP partner.

Result

You can reconfigure any ARS endpoint or the source ID value for SAML 1.x if you are editing an existing connection.

Configuring OpenID Provider information

You must configure OpenID Provider (OP) settings and information when configuring service provider (SP) browser single sign-on (SSO).

Steps

1. On the **OpenID Provider Info** tab, provide the scopes, the endpoints, and the authentication scheme. See the following image for an example.





If you have chosen to load the metadata from the OpenID Provider (OP) on the General Info tab, the **Scopes** field and all endpoints are pre-populated, provided that the metadata contains the information.

Field	Description
Scopes	The scopes to be included in the OpenID Connect authentication and OAuth token requests to the OP. Multiple space-separated values are allowed.
	The default value, without loading metadata from the OP, is openid.
	Tip:
	For a list of OpenID Connect defined scopes, see the section about requesting claims using scope values in the OpenID Connection specification at openid.net/specs/openid-connect-core-1_0.html#ScopeClaims.
Authorization	The authorization endpoint at the OP.
Endpoint	You can enter a relative path, starting with a forward slash, if you provide base URL on the General Info tab.
	There is no default value without loading metadata from the OP.

Field	Description
OpenID Connect Login Type	The OpenID Connect client profile of the client. This client represents PingFederate and is created and managed at the OP.
	 If the client is configured to support the Basic Client profile, select Code.
	The resulting value of the response_type parameter is code. If the client is configured to support the Implicit Client profile, select Form POST.
	The resulting value of the response_type parameter is id_token. If the client is configured to support the Implicit Client profile and the target application requires the associated access token, select Form POST with access token.
	The resulting values of the response_type parameter are id_token token.
	The default selection, without loading metadata from the OP, is Code.
Authentication Scheme	The client authentication method that PingFederate uses. Applicable and visible only to clients supporting the Basic Client profile.
	 Select Basic to submit credentials with HTTP Basic authentication. Select POST to submit credentials with POST. Select Private Key JWT to authenticate with the private_key_jwt Client Authentication method, see Client Authentication in the OpenID Connect specification.
	The default selection, without loading metadata from the OP, is Basic .

Field

Description

Authentication Signing Algorithm

Select the algorithm that PingFederate uses to sign the JWT.

Applicable and visible only when **Private Key JWT** is the chosen authentication scheme.

If PingFederate is either deployed to run in a Java 11 runtime environment or integrated with a hardware security module (HSM) and configured to use static keys for OAuth and OpenID Connect, additional RSASSA-PSS signing algorithms become available for selection. For more information on HSM integration and static keys, see and .



If static keys for OAuth and OpenID Connect are enabled, Elliptic-curve cryptography (EC) algorithms that have not been configured with an active static keys are hidden.

Changes made in the static-key configuration might affect runtime transactions and require additional changes here. For more information, see.



Note:

Based on the chosen signing algorithm, PingFederate selects the signing JSON Web Key (JWK) from its JWK Set (JWKS) at runtime.

In order for the OP to validate the signed JWT, ensure that the OP can access your PingFederate JWKS endpoint, which returns the current set of JSON Web Keys. The PingFederate JWKS endpoint is located at <Base URL>/pf/JWKS, where Base URL is defined on System # Server # Protocol Settings # Federation Info.

For example, if the Base URL field value is https://www.example.com, the PingFederate JWKS endpoint is https://www.example.com/pf/JWKS. You can pass the PingFederate JWKS endpoint directly to the OP or have the OP contact the PingFederate OpenID Provider configuration endpoint to obtain the information. For more information, see .

Field	Description
Token Endpoint, UserInfo Endpoint,	Various OAuth 2.0 and OpenID Connect 1.0 endpoints at the OP. For more information, see openid.net/connect.
and JWKS URL	Token Endpoint
	The Token Endpoint field is only visible and required for clients supporting the Basic Client profile. (In other words, the OpenID Connect Login Type field is set to Code .)
	UserInfo Endpoint
	The UserInfo Endpoint field is optional. If omitted, PingFederate only has access to the end-user claims from the ID tokens.
	JWKS URL
	The JWKS URL is required in order for PingFederate to validate the inbound ID tokens from the OP. If the OP signs its JWTs using an RSASSA-PSS signing algorithm, PingFederate must be deployed to run in a Java 11 runtime environment or integrated with a hardware security module (HSM) to process the digital signatures. For more information on HSM integration, see .
	There are no default values without loading metadata from the OP.
Sign Request	Select this check box to send request parameters as claims in a request object, a self-contained, signed JWT as one request query parameter to the OP.
	When this optional configuration is enabled, the OP can validate the integrity of the request parameters based on the digital signature found in the signed JWT. For more information, see the section explaining passing a request object by value in the OpenID Connect specification at openid.net/specs/openid-connect-core-1_0.html#RequestObject.
	This check box is not selected by default, in which case PingFederate sends request parameters with multiple query parameters, unsigned.

Field

Description

Request Signing Algorithm

Select the algorithm that PingFederate uses to sign the request object.

Applicable and visible only when the **Sign Request** check box is selected.

If PingFederate is either deployed to run in a Java 11 runtime environment or integrated with a hardware security module (HSM) and configured to use static keys for OAuth and OpenID Connect, additional RSASSA-PSS signing algorithms become available for selection. For more information on HSM integration and static keys, see and .



Note:

If static keys for OAuth and OpenID Connect are enabled, Elliptic-curve cryptography (EC) algorithms that have not been configured with an active static keys are hidden.

Changes made in the static-key configuration might affect runtime transactions and require additional changes here. For more information, see.



Note:

PingFederate automatically selects the signing JSON web key (JWK) based on the selected signing algorithm from its JWK Set (JWKS).

In order for the OP to validate the signed request object, ensure that the OP can access your PingFederate's JWKS URL, which returns the current set of JSON web keys. The PingFederate JWKS URL is located at <Base URL>/pf/JWKS, where Base URL is defined on System # Server # **Protocol Settings # Federation Info.**

For example, if the **Base URL** field value is https://www.example.com, the PingFederate JWKS URL is https://www.example.com/pf/JWKS. You can pass the JWKS URL directly to the OP or have the OP contact the PingFederate OpenID Provider configuration endpoint for it. For more information, see .

2. Optional: Remain on the OpenID Provider Info tab and specify the request parameters that are allowed to be included in the authentication requests to the OP under Request Parameters. For more information, see .

Configuring default target URLs

You can define default target URLs for identity provider (IdP) connections.

About this task

Use the Protocol Settings # Overrides tab to assign a default target URL for this IdP connection.

Steps

Optional: Enter a URL in the **Default Target URL** field.

If specified, the value overrides the default target setting defined in the Applications # Integration # SP Default URLs window.



Note:

The SAML 1.x specifications for IdP-initiated single sign-on (SSO) require a target URL to be specified. If the target application is specified in the URL parameter to the /sp/startSSO.ping application endpoint, any URL specified in the **Default Target URL** field in this window will not be used for those transactions.

Overriding authentication context in an IdP connection

You can map authentication context values between the local and remote values in an OpenID Connect or a SAML 2.0 identity provider (IdP) connection.

About this task

This optional configuration overrides how authentication context values are communicated with partners in both the authentication or authorization requests and their responses. Any values that are not defined in this configuration are passed through as-is.

As needed, you can use an asterisk, *, to match any values, a blank value for a scenario where the partner or the local request does not specify an authentication value, or both.

Steps

- 1. Go to Authentication # Integration # IdP Connections.
- 2. Click the name of the connection to open it in the IdP Connection window.
- On the Activation & Summary tab, scroll down to the Protocol Settings section, then click Overrides.
- 4. On the **Overrides** tab, specify the **Local** and **Remote** entry, then click **Add**.
- 5. Repeat the previous step to define additional mappings.
 - Click **Edit**, **Update**, or **Cancel** to make or undo a change to an existing entry. Click **Delete** or **Undelete** to remove an existing entry or cancel the removal request.
- 6. Click **Save** to complete the configuration.

Alternatively, click **Next** to carry on with the rest of the connection settings.

Example

Example

Suppose you are the service provider (SP) and your target application requires either the urn:oasis:names:tc:SAML:2.0:ac:classes:Kerberos or urn:oasis:names:tc:SAML:2.0:ac:classes:unspecified authentication context. While the IdP is capable of authenticating its users using a Kerberos-based authentication system, a proprietary identity

is capable of authenticating its users using a Kerberos-based authentication system, a proprietary identity management system, and a few internal web portals, the authentication context values are different than what your application supports. The authentication context values from the IdP are as follows.

Authentication method	AuthnContext values	
Kerberos-based authentication system	KerberosAuth	
Internal web portals	password, portal, or web	
Proprietary identity management system	No authentication context information is provided	

To override the AuthnContext values from the IdP, you can configure the IdP connection with the following authentication context mappings.

Local	Remote
urn:oasis:names:tc:SAML:2.0:ac:classes	: Kentberros Auth
urn:oasis:names:tc:SAML:2.0:ac:classes	:umspecified

Local	Remote	
urn:oasis:names:tc:	SAML:2.0:ac:classes:unspecified	

The first entry maps KerberosAuth to urn:oasis:names:tc:SAML:2.0:ac:classes:Kerberos. The second entry maps any authentication context values including password and portal to urn:oasis:names:tc:SAML:2.0:ac:classes:unspecified. The last entry overrides the authentication value to urn:oasis:names:tc:SAML:2.0:ac:classes:unspecified in the event that the assertion does not contain any authentication context information.

Configuring signature policy

The Signature Policy tab provides options controlling how digital signatures are used for SAML and WS-Federation single sign-on (SSO) messages.

About this task

The choices made on this tab depend on your partner agreement. For more information, see *Digital signing* policy coordination on page 88.

Digital signing is required for SAML response messages sent from the identity provider (IdP) through POST or redirect for SAML 2.0. The SAML specifications allow the signing of the entire SAML response message or the assertion portion inside the SAML response message. If you and your partner agree on the latter, select the Specify additional signature requirements and Require signed SAML Assertions options on this tab. When the latter is selected, only the assertion portion of the SAML response message is signed, not the entire SAML response message. This is the only option that appears for SAML 1.x and WS-Federation connections.

SAML 2.0 authentication requests from the service provider (SP) can also be signed to enforce security. This option appears only for SAML 2.0 connections and when the SP-initiated SSO profile is enabled on the SAML Profiles tab.

Steps

To continue, select the options based on your partner agreement.

Result

If you are editing an existing connection, you can reconfigure the digital signature policy, which might require additional configuration changes in subsequent tasks.

Specifying XML encryption policy (for SAML 2.0)

For SAML 2.0 configurations, in addition to using signed assertions to ensure authenticity, you and your partner can also agree to encrypt all or part of an assertion to improve privacy.

About this task

You can configure these settings on the **Encryption Policy** tab.



Note:

For WS-Fed connections with SAML 2.0 assertions, you cannot encrypt the entire assertion.

Option	Name identifier (SAML_SUBJE		Encrypt the SAML_SUBJECT in SLO messages to the IdP	Allow encrypted SAML_SUBJECT in SLO messages from the IdP
None	No encryption.	No encryption.	No encryption.	No encryption.

Option	Name identifier (SAML_SUBJEC		Encrypt the SAML_SUBJECT in SLO messages to the IdP	Allow encrypted SAML_SUBJECT in SLO messages from the IdP
The entire assertion	Encryption allowed.	Encryption allowed.	Encryption allowed as an available option.	Encryption allowed as an available option.
SAML_SUBJECT (Name Identifier)	TEncryption allowed.	Encryption allowed as an available option.	Encryption allowed as an available option.	Encryption allowed as an available option.
One or more attributes	Encryption allowed.	Encryption allowed as an available option.	Encryption allowed as an available option only if you select to allow the entire assertion or the SAML_SUBJECT to be encrypted.	Encryption allowed as an available option only if you select to allow the entire assertion or the SAML_SUBJECT to be encrypted.

To enable encryption:

Steps

- 1. Click the Allow encrypted SAML Assertions and SLO messages option.
- 2. Choose whether this identity provider (IdP) partner will encrypt the entire assertion, the SAML SUBJECT name identifier, one or more other attributes, or some combination.
- 3. If your partner is encrypting the name identifier, indicate whether you will encrypt this attribute in outbound SAML 2.0 single logout (SLO) messages, allow its encryption for inbound messages, or both.

Result

If you are editing an existing connection, you can reconfigure the XML encryption policy, which might require additional configuration changes in subsequent tasks.

Reviewing protocol settings for SP browser SSO

Confirm your protocol settings for browser-based single sign-on (SSO) before completing configuration.

Steps

- To amend your configuration, click the corresponding tab title, then follow the configuration wizard to complete the task.
- To keep your changes, click **Done** and continue with the rest of the configuration.



When editing an existing configuration, you can also click **Save** as soon as the administrative console offers the opportunity to do so.

To discard your changes, click Cancel.

Reviewing Browser SSO settings

You can review your browser single sign-on (SSO) settings before completing configuration.

Steps

 To amend your configuration, click the corresponding tab title, then follow the configuration wizard to complete the task.



When editing an existing configuration, you can also click **Save** as soon as the administrative console offers the opportunity to do so.

To discard your changes, click Cancel.

Manage the Attribute Query profile in an IdP connection

At the attribute query step, you configure your connection to request user attributes from your partner identity provider (IdP), if you chose this option.

For more information on IdP connection options, see Choosing IdP connection options on page 665. Attribute queries do not depend on single sign-on (SSO), but can be used independently or in conjunction with browser SSO or provisioning to provide flexibility in how a user authenticates with service provider (SP) applications. For more information, see Attribute Query and XASP on page 43 and the /sp/ startAttributeQuery.ping on page 1142 SP application endpoint.

Setting the Attribute Authority Service URL

The attribute authority service URL corresponds to the endpoint location where your identity provider (IdP) provider receives attribute query requests.

About this task

Attribute authority is the term used to refer to an IdP that provides user attributes to an attribute requester or your service provider (SP) site. For more information, see Attribute Query and XASP on page 43.

Steps

 Enter the fully qualified URL or a relative path if you have defined a base URL on the General Info tab. For more information, see *Identifying the partner* on page 668.

Mapping attribute names for Attribute Query

If the application at your site uses different names for user attributes than the names defined by the attribute authority, you must map them on the Attribute Name Mapping tab.

About this task

When the service provider (SP) receives a request from a local application to send an attribute query to this attribute authority partner, the requested user attributes are replaced with the names mapped here.

You must predetermine this information in your agreement with this connection partner.

Steps

To map an attribute, configure the Local Name and Remote Name fields for the attribute, then click Add.

Choice	Action
Modify an attribute name	Click Edit under Action for the attribute, make the change, then click Update .
	Note:
	If you change your mind, ensure that you click Cancel under Actions and not the Cancel button, which discards any other changes you might have made in the configuration steps.

Choice	Action
Delete an attribute	Click Delete under Action for the attribute.

Configuring security policy for Attribute Query

The **Security Policy** tab allows you to specify the digital signing and encryption policy to which you and your partner have agreed.

About this task

These selections will trigger requirements for setting up credentials. For more information, see *Configuring* security credentials on page 731.

This tab also allows you to mask incoming attribute values in log files. For more information, see Attribute masking on page 95. When you enable this selection, all user attributes returned from this identity provider (IdP) are masked.

Steps

Select or clear the check boxes the relevant check boxes.

Reviewing the Attribute Query settings

You can review your attribute query settings before completing configuration.

Steps

- To amend your configuration, click the corresponding tab title, then follow the configuration wizard to complete the task.
- To keep your changes, click **Done** and continue with the rest of the configuration.



Tip:

When editing an existing configuration, you can also click **Save** as soon as the administrative console offers the opportunity to do so.

To discard your changes, click Cancel.

Configuring just-in-time provisioning

PingFederate's just-in-time (JIT) provisioning allows service providers (SPs) to create user accounts on the fly during single sign-on (SSO) events, based on attributes received in SSO tokens from identity providers (IdPs).

About this task

An SP can also use JIT provisioning to update existing user records.



This configuration task is presented in the administrative console only when the JIT Provisioning check box is selected on the Connection Options tab.



- 1. Go to Authentication # Integration # IdP Connections.
- 2. Create a new IdP connection or select an existing IdP connection .
- 3. On the Connection Type tab, select the Browser SSO Profiles check box and a protocol from the
- 4. On the Connection Options tab, select the Browser SSO check box and then the JIT Provisioning check box.
- 5. Complete the **Browser SSO** configuration.
- 6. On the JIT Provisioning tab, click Configure User Provisioning to begin the configuration of JIT provisioning.

Selecting attribute sources (SAML 2.0)

For SAML 2.0 connections, the server can be configured to use only assertion attributes for user provisioning, or to retrieve more attributes from the identity provider (IdP) in a follow-on attribute query transaction.

About this task

The **User Attributes** tab displays the attributes expected in the assertion from this IdP.





Note:

The attribute query is a SAML 2.0 profile. For OpenID Connect, SAML 1.x, and WS-Federation connections, this tab is not presented. PingFederate uses only attributes from the assertion for user provisioning.

Steps

 If you and your IdP partner have agreed to use the Attribute Query profile for provisioning, select that option before leaving this tab.

You configure the attribute query profile later in the task flow.

Identifying the user repository

PingFederate's just-in-time (JIT) provisioning supports several directory servers and Microsoft SQL Server.

About this task

For more information on versions and requirements supported by PingFederate, see System requirements on page 110.



Note:

PingFederate was tested with vendor-specific JDBC drivers. For more information, see *Database driver* information on page 115.

On the User Repository tab, select a datastore from the Active Data Store list.





If the desired datastore is not shown in the list, PingFederate has not been configured to access it. Click **Manage Data Stores** to configure your datastore.

- If you are using an LDAP store, see the sections immediately following:
 - Specifying an LDAP user-record location on page 701
 - Entering an LDAP filter on page 701
 - Identifying provisioning attributes for LDAP on page 702
- If you are using a Microsoft SQL Server, skip to this section:
 - Choose a SQL method

Specifying an LDAP user-record location

After choosing a datastore, indicate where in the store PingFederate should write new user records or update existing ones.

About this task





Note:

The Location tab appears only when an LDAP datastore is chosen on the User Repository tab.

Steps

Enter the base distinguished name (DN) in the Base DN field.

A base DN is the DN of the tree structure in which the search begins. Leave this field blank if records are located at the LDAP root.

Entering an LDAP filter

On the Unique User ID tab, create an LDAP filter to identify user accounts to be provisioned or updated during single sign-on (SSO) events.

About this task

PingFederate uses this expression in conjunction with the Base DN value defined on the Location tab to locate existing account records and to add new ones.





This tab appears only when an LDAP datastore is chosen on the User Repository tab.

Steps

Enter the statement in the Filter field.

The filter is in the form: attribute=\${value}.



Unlike filters used to retrieve LDAP attributes for adapter mapping, do not enclose the statement in parentheses.

The left-side variable is an attribute in your user-datastore. Click the link near the lower-left corner of the tab to see a list of available attributes.

The right side of the filter uses one or more attribute values passed in from the SSO token. Variables for these attributes, including the correct syntax, are listed under JIT Attributes.



If you are unfamiliar with writing LDAP queries, see the documentation accompanying your LDAP installation.

Identifying provisioning attributes for LDAP

On the Attributes tab, you can select the datastore attributes to be provisioned.

About this task





This tab appears only when an LDAP datastore is chosen on the **User Repository** tab.

Steps

- 1. Select a root object class and an attribute from the lists, and then click **Add Attribute**.
- 2. Repeat for each attribute requiring provisioning.

PingFederate allows you to map attributes directly to a single database table, the default, or to SQL storedprocedure parameters for Java Database Connectivity (JDBC) datastores,

About this task

Choose and configure the preferred method on the **SQL Method** tab.





This tab appears only when you specify Microsoft SQL Server datastore on the User Repository tab.

Steps

Make a selection as needed and click Next.

Depending on the selection, different steps appear under the JIT Provisioning task. See the sections indicated for more information.

- If mapping attributes directly to a table, see the topics sections immediately following:
 - Specifying a database user-record location on page 703
 - Specifying a unique ID database column on page 704
- If using a stored procedure, skip to Specifying a stored procedure location on page 704

Specifying a database user-record location

For database provisioning to a table, you must indicate where PingFederate should write new user records or update existing ones.

About this task

Configure the location settings on the **Location** tab.





Note:

This tab appears only when you choose a Microsoft SQL Server datastore on the User Repository tab and select the Table option on the SQL Method tab.

Steps

On the Location tab, select the database schema and the table.

Field	Description	
Schema	Select the table structures that store information within the database.	
Table	Select the name of the database table that contains user records.	

Field	Description
Columns to fulfill	For the selected table, all attributes and their data types are displayed.
	If the list of columns is not current due to any recent changes, click Refresh .
	On the Attribute Fulfillment tab, all attributes must be mapped for the database insertion to succeed, although a null entry can be used for optional attributes.



Click the link near the lower-left corner of the tab to see a list of available attributes from the single sign-on (SSO) token.

Specifying a unique ID database column

PingFederate uses the database column you specify on the Unique ID tab to check whether a user record already exists for the incoming single sign-on (SSO) token.

About this task





Note:

This tab only appears if you have chosen a Microsoft SQL Server datastore on the User Repository tab and you selected the Table option on the SQL Method tab.

Steps

Select a column that represents a unique characteristic about the database entry for a particular user.

Specifying a stored procedure location

If you are using a stored procedure for provisioning the user database, specify its location on the Stored Procedure Location tab.

About this task





Note:

This window appears only when you have chosen a Microsoft SQL Server datastore on the User Repository tab and the you have selected the Stored Procedure option on the SQL Method tab.

- 1. From the **Schema** menu, select the table structure that contains stored procedure within the database.
- 2. From the **Stored Procedure** menu, select the stored procedure needed to provision the user database.



Important:

The database account used by PingFederate must have access to the schema in which the stored procedure is located and execute permission for the procedure.

Result: For the selected procedure, all parameters and their data types display in the **Procedure** parameters to fulfill section.



Note:

To see a list of available attributes from the single sign-on (SSO) token, click View Attribute Contract.

You map attribute values from the SSO token to the parameters on the Attribute Fulfillment tab.

Troubleshooting:

If the list of parameters is not current due to any recent procedure revisions, click Refresh.

Mapping attributes to a user account

Map incoming attributes to the account attributes on an LDAP server, the columns in a database table on a Microsoft SQL Server, or the parameters of a Microsoft SQL Server stored procedure.

About this task

In addition to values obtained from the single sign-on (SSO) token, you can map attributes from the context of the SSO token text, with or without reference values from the SSO token, and expression if enabled.

If you select a Microsoft SQL Server datastore on the User Repository tab, then on the Attribute Fulfillment tab you can test the insertion of attribute values into the database table or the stored procedure. When mapping to a database column of the datetime or smalldatetime data type, if you are not using a stored procedure to convert the incoming string value, you can use a PingFederate Java conversion method through OGNL expressions.

- 1. On the Attribute Fulfillment tab, select a source from the list for each target attribute or parameter. Choose from:
 - Assertion or Provider Claims

Values are contained in the SSO token from this identity provider (IdP). When you select this, the associated **Value** list is populated by the attribute contract.

Context

Values are returned from the context of the transaction at runtime.



Note:

As the HTTP Request is retrieved as a Java object rather than text, OGNL expressions are more appropriate to evaluate and return values. Choose Expression from the list and then click Edit to enter an expression.

Attribute Query

This choice appears only if you choose the **Attribute Query** profile for provisioning.

To map an attribute-query value, use the syntax \${query attribute}. You can combine attribute-query values with references to attributes in the attribute contract; for example, \${query attribute}+\${attribute.

References to attributes not contained in the attribute contract result in an attribute query back to the IdP partner.

Expression



Tip:

Enable OGNL expression by editing the <pf install>/pingfederate/server/default/ data/config-store/org.sourceid.common.ExpressionManager.xml file. Restart PingFederate after saving the change.

For a clustered PingFederate environment, edit the

org.sourceid.common.ExpressionManager.xml file on the console node, sign on to the

administrative console to replicate this change to all engine nodes in the System # Server # Cluster Management window, and restart all nodes.

This option provides more complex mapping capabilities, such as transforming incoming values into different formats. All of the variables available for text entries are also available for expressions.



Tip:

If you need to map multiple attribute values from one or multiple sources to one attribute value, use an OGNL expression to create it.

For database mapping, if the data type of a target parameter is datetime or smalldatetime, you can use an expression to convert date-time strings from the SSO token. After selecting Expression, click Datetime OGNL Examples for syntax information and examples.

System Managed

This mapping option appears only when any automatically assigned attributes are among columns to be provisioned, such as an identity or a timestamp column on the Microsoft SQL Server.

Text

The value is what you enter. This can be text only, or you can mix text with references to any of the values from the SSO token, using the \${attribute} syntax.



Note:

For LDAP mapping, choose **Text** as the **Source** for the objectClass attribute.

For mapping into a database, if no entry is required for a column, you can leave the field blank. A blank entry results in an empty string in the database for string data types and null for all other data types. Alternatively, for string types, you can enter null in the field to explicitly set null in the column.

2. Select or enter an attribute value.

All values must be mapped. For optional table columns, you can leave the field blank or, for string data types, enter null to avoid empty strings.

No value is required for **System Managed** attributes.



Note:

For Active Directory, enter user in the objectClass field. For Oracle Directory Server or Oracle Unified Directory, enter inetOrgPerson.

- 3. Optional: When mapping to a Microsoft SQL Server datastore, test the insertion. Choose from:
 - If testing from a table:
 - a. Click Test insert into .
 - b. Enter values for each applicable target parameter.
 - c. Click Test Insert.

If the test succeeds, a confirmation displays along with the values inserted.



CAUTION:

Unless you want to keep the test values in the database, click Roll Back All Test Inserts.

- If testing from a stored procedure:
 - a. Click Test call to cedure>.
 - b. Enter values for each applicable target parameter.
 - c. Click Test Stored Procedure Call.

For stored procedures, only a confirmation displays if the test is successful, indicating that the procedure was populated with parameter values.



CAUTION:

No roll back feature is provided because PingFederate does not know the result of the procedure. Database rollback must be handled manually.

When finished, click Return to Attribute Fulfillment.

Choosing an event trigger

Choose whether PingFederate initiates user provisioning only when the user identifier is new, or every time your site receives a single sign-on (SSO) token.

About this task

If you choose to have PingFederate initiate user provisioning every time your site receives an SSO token, for all SSO tokens, an existing user account is always updated with incoming attributes.





Note:

This tab does not appear for a Microsoft SQL Server datastore if provisioning is accomplished using a stored procedure, because the procedure is always called for all SSO tokens. The procedure should handle both provisioning new users and updating existing ones.

- On the Event Trigger tab, in the Specify the trigger that initiates a user-provisioning event section, select one of the following: Choose from:

 - Only SAML Assertations Containing a New User ID
 - All SAML Assertations

Configuring an error handling method

If user provisioning fails for any reason during single sign-on (SSO) events, you can choose to continue the process by passing the user's attributes to your target application or to abort the SSO transaction.

About this task

When SSO is aborted, the user is redirected to an error page and the failure is written to the server log.



Steps

- On the Error Handling tab, in the How should the server handle a failure in a user provisioning request? section, select one of the following: Choose from:
 - Send the User's Attributes to the Target Application
 - Abort the SSO Transaction

Reviewing the JIT provisioning configuration

You can review your just-in-time (JIT) provisioning settings before completing configuration.

Steps

- To amend your configuration, click the corresponding tab title, then follow the configuration wizard to complete the task.
- To keep your changes, click **Done** and continue with the rest of the configuration.



When editing an existing configuration, you can also click Save as soon as the administrative console offers the opportunity to do so.

To discard your changes, click Cancel.

Configuring SCIM inbound provisioning

The System for Cross-domain Identity Management (SCIM) inbound provisioning configuration provides for a two-way mapping of attributes.

About this task

The first way facilitates SCIM operations used to create and update records in the datastore. The second way allows the same SCIM client to retrieve those records and have the attribute values mapped back to their corresponding designation in the client store. The dual mapping provides greater flexibility, especially for OGNL-expression transformations (for example, converting two attributes into one multivalued attribute and then back again). For more information, see Writing user information to the datastore on page 716 and Configuring a SCIM response on page 719.



Note:

SCIM-client requests must include authentication credentials, which you configure on the Credentials # Back-Channel Authentication tab. The same credentials needed for single sign-on (SSO), are also used for SCIM transactions.

Steps

- 1. On the Authentication # Integration # IdP Connections window, create a new IdP connection or select an existing IdP connection.
- 2. On the Connection Type tab, select the Inbound Provisioning check box and one of the following options:

Choose from:

- User Support
- User and Group Support
- 3. On the Inbound Provisioning tab, click Configure Inbound Provisioning to begin the configuration of SCIM inbound provisioning.



Specifying the user repository

PingFederate supports Active Directory (AD) user stores and custom identity store provisioners for inbound provisioning.

About this task

Manage the datastore serving as the local repository on the **Repository** tab.



Active Directory user stores require an LDAPS connection to the datastore.

- Choose one of the following options, and then specify the datastore from the list. Choose from:
 - Active Directory Data Store
 - Identity Store Provisioner





If the correct datastore is not shown in the list, then PingFederate is not configured to access the store. Click Manage Data Stores to set up the desired datastore.



Identifying an LDAP user-record location

After choosing a datastore, you can indicate where in the datastore user and group records exist so PingFederate can create, read, update, or delete or disable them.

About this task

These settings can be configured on the **Location** tab.





Note:

This tab appears only if you are configuring an LDAP user store for provisioning.

Steps

 Enter the base distinguished name (DN) of the tree structure where user records are stored in the Base DN field.



Note:

PingFederate looks only at this node level or below it for user accounts that need provisioning

Defining a unique user ID

On the Unique user ID tab, you can create an LDAP filter to resolve user accounts for System for Crossdomain Identity Management (SCIM) operations.

About this task

PingFederate uses LDAP filter in conjunction with the Base DN value, defined on the Location tab, to add new account records.





This tab only appears if you are configuring an LDAP user store for provisioning.

Steps

Enter the statement in the Filter text field.

The filter is in the form: attribute=\${value} where attribute is an attribute in your userdatastore and value is the attribute value or values passed in from the SCIM request. To see a list of available attributes in your user-datastore, click View List of Available LDAP Attributes. Variables for these attributes, including the correct syntax, are listed under SCIM Attributes.



Note:

Unlike filters used to retrieve LDAP attributes for adapter mapping, do not enclose the statement in parentheses.



If you are unfamiliar with writing LDAP queries, see the documentation accompanying your LDAP installation.

Defining a unique group ID

On the Unique Group ID tab, you can create an LDAP filter to resolve groups for System for Crossdomain Identity Management (SCIM) operations.

About this task

PingFederate uses this LDAP filter in conjunction with the **Base DN** value, defined on the **Location** tab, to add new groups.





This tab appears only if you are configuring an LDAP user store for provisioning and you have selected the User and Group Support option on the Connection Type tab.

Steps

Enter the statement in the Filter text field.

The filter is in the form: attribute=\${value} where attribute is an attribute in your userdatastore and value is the attribute value or values passed in from the SCIM request. To see a list of available attributes in your user-datastore, click View List of Available LDAP Attributes. Variables for these attributes, including the correct syntax, are listed under **SCIM Attributes**.



Note:

Unlike filters used to retrieve LDAP attributes for adapter mapping, do not enclose the statement in parentheses.



If you are unfamiliar with writing LDAP queries, see the documentation accompanying your LDAP installation.

Defining custom SCIM attributes

When configuring System for Cross-domain Identity Management (SCIM) inbound provisioning, you can define custom attributes.

About this task

PingFederate supports SCIM attributes in the core schema and custom attributes through a schema extension.



Note:

Custom attributes are optional. If your use case does not require any additional attributes, click Next on the Custom SCIM Attributes tab.

To support custom attributes, you must specify the schema extension and the custom attributes in the connection. There are four attribute types:

- Simple attributes
- Simple multivalued attributes
- Complex attributes
- Complex multivalued attributes

The following fragment illustrates a SCIM message supporting schema extension urn:scim:schemas:extension:custom:1.0 with four attributes, one of each attribute type. The table afterward describes the details of each attribute.

```
"userName": "CBrown",
"active":true,
"schemas":[
 "urn:scim:schemas:core:1.0",
```

```
"urn:scim:schemas:extension:custom:1.0"
],
"urn:scim:schemas:extension:custom:1.0":{
   "supervisor": "JSmith",
  "territories":[
     "Montana",
     "Idaho",
     "Wyoming"
  ],
   "options":{
     "quantity":"10000",
"strike" :"5.25",
"first" :"2017-12-01",
"last" :"2025-03-31"
   "tablets":[
     {
        "model" :"8086",
        "serial":"5500-2020-965",
"type" :"office"
        "model" : "8088",
        "serial":"5500-2040-151",
"type" :"remote"
  ]
}
```

Attribute Name	Attribute Type	Sub-Attributes (Complex)
supervisor	Simple	Not applicable
territories	Simple multivalued	Not applicable
options	Complex	quantity, strike, first, and last
tablets	Complex multivalued	model, serial, and type.
		Note: type is a reserved sub-attribute for a complex multivalued attribute.



For more information about SCIM and attribute types, see the website www.simplecloud.info.

Steps

Go to the Custom SCIM Attributes tab.



Choose from:

 To specify a schema extension, enter the URI of the schema extension in the Extension Namespace field



The default value is urn: scim: schemas: extension: custom: 1.0. You can keep this value if your partner identifies custom attributes by this URI in its SCIM messages.

- To add a custom attribute, enter an attribute name and click Add. Repeat this step to add more custom attributes as needed.
- To delete a custom attribute, click **Delete** next to the custom attribute.
- To undo the deletion, click Undelete.
- To edit a custom attribute, click **Edit** next to the custom attribute.

Result:

The administrative console displays the **Custom SCIM Attribute Options** tab, where you can:

- Change the attribute name.
- Set the attribute as a simple multivalued attribute.
- Add sub-attributes to make it a complex attribute.
- Add sub-attributes and set the attribute as a complex multivalued attribute.

For more information, see Configuring custom SCIM attribute options on page 715.

Configuring custom SCIM attribute options

After choosing a custom System for Cross-domain Identity Management (SCIM) attribute on the Custom SCIM Attribute tab, you can use the Custom SCIM Attribute Options tab to configure the attribute.

About this task

On this tab you can change the attribute name, set the attribute as a simple multivalued attribute, add sub-attributes to make it a complex attribute, and add sub-attributes and set the attribute as a complex multivalued attribute.

Steps

Go to the Custom SCIM Attribute Options tab.

Choose from:

To change the name of the custom attribute, replace the current value in the Name field, and then click



Done.

To define the custom attribute as a simple multivalued attribute, select the Is Multi-valued check box,



and then click Done.

To define the custom attribute as a complex attribute, enter a sub-attribute, and then click Add.
 Repeat this step to add more sub-attributes as needed, and then click Done.



Click **Edit**, **Update**, or **Cancel** to make or undo a change to the name of a sub-attribute. Click **Delete** or **Undelete** to remove a sub-attribute or cancel the deletion.



- To define the custom attribute as a complex multivalued attribute, follow these steps.
 - a. Enter a sub-attribute, and then click Add. Repeat this step to add more sub-attributes as needed.



Click **Edit**, **Update**, or **Cancel** to make or undo a change to the name of a sub-attribute. Click **Delete** or **Undelete** to remove a sub-attribute or cancel the deletion.

- b. Select the Is Multi-valued check box.
- c. If you have chosen Active Directory as your user store, you must specify at least one value in the **Types** section for type, a reserved sub-attribute for a complex multivalued attribute. For more information, see *Specifying the user repository* on page 710



Writing user information to the datastore

To configure how PingFederate completes create and update operations for user accounts from a System for Cross-domain Identity Management (SCIM) request, identify incoming attributes and map them to datastore attributes.

Steps

On the Write Users tab, click Configure Write Users to continue.



Identifying inbound provisioning attributes for LDAP

When configuring System for Cross-domain Identity Management (SCIM) inbound provisioning, you must identity the attributes you want to provision.

About this task

You can select the datastore attributes you want to provision on the Attributes tab.



This tab appears only if you are configuring an LDAP user store for provisioning.



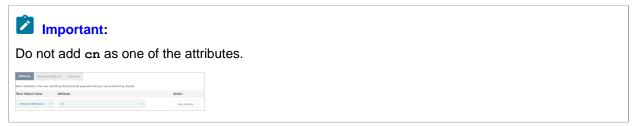
The following attributes are managed internally by PingFederate and do not require mapping:

- objectClass
- unicodePwd
- objectGUID
- userAccountControl

You can override the internal management of objectClass and unicodePwd by selecting these attributes and mapping them to SCIM attributes on the Attribute Fulfillment tab. In this case, the values you supply are used. The objectGUID and userAccountControl attributes cannot be overridden and are ignored if selected.

Steps

1. Select a root object class and an attribute from the lists, then click **Add Attribute**.



2. Repeat the previous step for each attribute requiring provisioning.

Mapping attributes to user accounts

Map attribute values in the System for Cross-domain Identity Management (SCIM) request to user-account attributes.

About this task



1. On the Attribute Fulfillment tab, for each attribute, select a source from the Source list and then choose or enter a value. You must map all attributes.

Context

When selected, the Value list populates with the available context of the transaction. Select the desired context from the list.



Note:

As the HTTP Request context value is retrieved as a Java object rather than text, use OGNL expressions to evaluate and return values.



Note:

If you are configuring an **OAuth Attribute Mapping** configuration and have added PERSISTENT GRANT LIFETIME as an extended attribute in the Authorization Server Settings window, you can set the lifetime of persistent grants based on the outcome of attribute mapping expressions or the per-client Persistent Grants Max Lifetime setting.

- To set lifetime based on the per-client Persistent Grants Max Lifetime setting, select Context from the Source list and Default Persistent Grant Lifetime from the Value list.
- To set lifetime based on the outcome of attribute mapping expressions, select **Expression** as the source and enter an OGNL expression in the Value field.

If the expression returns a positive integer, the value represents the lifetime of the persistent grant in minutes.

If the expression returns the integer 0, PingFederate does not store the grant and does not issue a refresh token.

If the expression returns any other value, PingFederate sets the lifetime of the persistent grant based on the per-client Persistent Grants Max Lifetime setting.

To set a static lifetime, select Text from the Source list and enter a static value in the Value field.

This is suitable for testing purposes, or cases where the persistent grant lifetime must always be set to a specific value.

Expression



Enable OGNL expression by editing the <pf install>/pingfederate/server/default/ data/config-store/org.sourceid.common.ExpressionManager.xml file. Restart PingFederate after saving the change.

For a clustered PingFederate environment, edit the

org.sourceid.common.ExpressionManager.xml file on the console node, sign on to the

administrative console to replicate this change to all engine nodes in the System # Server # Cluster Management window, and restart all nodes.

This option provides more complex mapping capabilities, such as transforming incoming values into different formats. All of the variables available for text entries are also available for expressions.



Tip:

If two attribute values from a SCIM request need to be mapped to one LDAP attribute value, use an OGNL expression to create it.

SCIM User

When you make this selection, the associated Value list populates with defined components of the SCIM request.

No Mapping

Select this option to ignore the **Value** field.

The value is what you enter. This can be text only, or you can mix text with references to any of the values from the SCIM request using the \${attribute}syntax.

2. Click Done.

Reviewing user mapping (Write Users) configuration

You can review your user mapping settings before completing configuration.

About this task

The **Summary** tab provides an overview of the inbound provisioning configuration for request mapping.

Steps

- To amend your configuration, click the corresponding tab title, then follow the configuration wizard to complete the task.
- To keep your changes, click **Done** and continue with the rest of the configuration.



When editing an existing configuration, you can also click Save as soon as the administrative console offers the opportunity to do so.

To discard your changes, click Cancel.

Configuring a SCIM response

To configure a System for Cross-domain Identity Management (SCIM) response to a request to read and return provisioned SCIM attributes, identify and map the user account attributes you want to include.

About this task

Begin this process on the **Read Users** tab.



On the Read Users tab, click Configure Read Users to continue.

Identifying expected user attributes for the SCIM response

An attribute contract is a set of user attributes that you and your partner agree will be sent in a System for Cross-domain Identity Management (SCIM) response for this connection.

About this task

On the **Attribute Contract** tab, the attributes you mapped to user account attributes on the **Write Users** tab appear under **Attribute Contract**.



Optionally, you can mask the values of attributes in the log files that PingFederate writes when it sends the SCIM response.

There are multiple SCIM attributes that are managed internally by PingFederate and are unavailable for inclusion in the attribute contract:

- id
- active

Steps

Click **Available SCIM Attributes** near the lower-left corner of the tab to include additional attributes you want to map in the SCIM response.

Option	Action
Add an attribute	Enter the attribute name in the text box, select the check box under Mask Values in Log as needed, then click Add .
	Attribute names are case-sensitive and must correspond to the attribute names expected by your partner. To see a list of available attributes, click Available SCIM attributes .
Modify an attribute name or masking selection	Click Edit under Action for the attribute, make the change, then click Update .
	Note: If you change your mind, make sure you click Cancel under Actions, not the Cancel button, which discards any other changes you might have made in the configuration steps.
Delete an attribute	Click Delete under Action for the attribute.

Identifying LDAP attributes for the SCIM response

On the Attributes tab, you can select the LDAP attributes you want to map to attributes in the System for Cross-domain Identity Management (SCIM) response.

About this task





This tab appears only if you are configuring an LDAP user store for provisioning.

Steps

- 1. Select a root object class and an attribute from the lists, then click **Add Attribute**.
- 2. Repeat for each attribute requiring provisioning.

Mapping attributes into the SCIM response

Map outgoing user-account attributes to System for Cross-domain Identity Management (SCIM) responses to READ requests.

Steps

- 1. On the Attribute Fulfillment tab, for each target attribute, select a source from the Source list, then choose or enter a value. All target attributes must be mapped. Choose from:
 - Context

When selected, the Value list populates with the available context of the transaction. Select the desired context from the list.



Note:

Because the HTTP Request context value is retrieved as a Java object rather than text, use OGNL expressions to evaluate and return values.



Note:

If you are configuring an OAuth Attribute Mapping configuration and have added PERSISTENT GRANT LIFETIME as an extended attribute in the Authorization Server Settings window, you can set the lifetime of persistent grants based on the outcome of attribute mapping expressions or the per-client Persistent Grants Max Lifetime Setting.

- To set lifetime based on the per-client Persistent Grants Max Lifetime setting, select Context from the Source list and Default Persistent Grant Lifetime from the Value list.
- To set lifetime based on the outcome of attribute mapping expressions, select **Expression** as the source and enter an OGNL expression in the Value field.

If the expression returns a positive integer, the value represents the lifetime of the persistent grant in minutes.

If the expression returns the integer 0, PingFederate does not store the grant and does not issue a refresh token.

If the expression returns any other value, PingFederate sets the lifetime of the persistent grant based on the per-client Persistent Grants Max Lifetime setting.

To set a static lifetime, select **Text** from the **Source** list and enter a static value in the **Value** field.

This is suitable for testing purposes, or cases where the persistent grant lifetime must always be set to a specific value.

Expression

This option provides more complex mapping capabilities, such as transforming outgoing values into different formats. All of the variables available for text entries are also available for expressions.



Tip:

If you need to map an LDAP attribute to two attributes in a SCIM response, use an OGNL expression to create them.



Enable OGNL expression by editing the <pf install>/pingfederate/server/default/ data/config-store/org.sourceid.common.ExpressionManager.xml file. Restart PingFederate after saving the change.

For a clustered PingFederate environment, edit the

org.sourceid.common.ExpressionManager.xml file on the console node, sign on to the administrative console to replicate this change to all engine nodes in the System # Server # Cluster Management window, and restart all nodes.

LDAP

Values are returned from your query. When you make this selection, the Value list populates with the LDAP attributes you identified for this datastore.

Identity Store

Values are returned from your query. When you make this selection, the Value list populates with the Identity Store attributes you identified for this datastore.

No Mapping

Select this option to ignore the Value field.

Text

The value is what you enter. This can be text only, or you can mix text with references to any of the values from the SCIM request, using the \${attribute} syntax.

Click Done.

Reviewing SCIM response (Read Users) configuration

You can review your System for Cross-domain Identity Management (SCIM) response mapping before completing configuration.

About this task

The Summary tab provides an overview of the inbound provisioning configuration for SCIM response mapping.

Steps

- To amend your configuration, click the corresponding tab title, then follow the configuration wizard to complete the task.
- To keep your changes, click **Done** and continue with the rest of the configuration.



Tip:

When editing an existing configuration, you can also click **Save** as soon as the administrative console offers the opportunity to do so.

• To discard your changes, click Cancel.

Configuring the handling of SCIM delete requests

You can use the Delete/Disable Users tab to define how System for Cross-domain Identity Management (SCIM) delete requests are handled within your user datastore.

About this task



Important:

If the group support option is enabled, when PingFederate receives a SCIM delete request for a group, it always removes the specified group from the datastore.





Note:

This tab appears only if you are configuring an LDAP user store for provisioning.

Steps

Click one of the two available options for **SCIM DELETE message behavior**.

Choose from:

 Click Disable User to make the user inactive within the datastore. This approach is preferred in situations where accounts must be retained for auditing reasons.

In order to be SCIM compliant when deleting users, PingFederate returns an HTTP 404 response code for all subsequent operations related to the user-effectively treating the user as if they have been deleted from the LDAP user store. For more information, see SCIM specifications.



CAUTION:

If the user is disabled through another method, PingFederate still treats that user as if they have been deleted and returns HTTP 404 response codes for all subsequent requests.

Click Permanently Delete User to remove the user from the datastore.

Writing group information to the datastore

To configure how PingFederate completes create and update operations for groups from a System for Cross-domain Identity Management (SCIM) request, identify incoming attributes and map them to datastore attributes.

About this task

You can identify and map these attributes on the **Write Groups** tab.



Steps

Click Configure Write Groups to continue.

Identifying inbound provisioning group attributes for LDAP

You must identify the datastore attributes you want to provision when writing group information to the datastore.

About this task

You can identify these attributes on the Attributes tab.





This tab only appears if you are configuring an LDAP user store for provisioning and the **User and Group Support** option is selected on the **Connection Type** tab.

PingFederate internally manages several attributes that do not require mapping:

- objectClass
- objectGUID
- member

You can override the internal management of objectClass by selecting and mapping it to a System for Cross-domain Identity Management (SCIM) attribute on the Attribute Fulfillment tab. In this case, the values you supply are used. The objectGUID and member attributes cannot be overridden and are ignored if selected.

Steps

- 1. Select a root object class and an attribute from the lists, and then click **Add Attribute**.
- 2. Repeat the previous step for each attribute requiring provisioning.

Mapping attributes to groups

Map attribute values in the System for Cross-domain Identity Management (SCIM) request to group attributes.

About this task

Steps

- 1. On the Attribute Fulfillment tab, for each attribute, select a source from the list and then choose or enter a value. You must map all target attributes.
 - Context

When selected, the Value list populates with the available context of the transaction. Select the desired context from the list.



Note:

As the HTTP Request context value is retrieved as a Java object rather than text, use OGNL expressions to evaluate and return values.



Note:

If you are configuring an OAuth Attribute Mapping configuration and have added PERSISTENT GRANT LIFETIME as an extended attribute in the Authorization Server Settings window, you can set the lifetime of persistent grants based on the outcome of attribute mapping expressions or the per-client Persistent Grants Max Lifetime Setting.

- To set lifetime based on the per-client Persistent Grants Max Lifetime setting, select Context from the Source list and Default Persistent Grant Lifetime from the Value list.
- To set lifetime based on the outcome of attribute mapping expressions, select **Expression** as the source and enter an OGNL expression in the Value field.

If the expression returns a positive integer, the value represents the lifetime of the persistent grant in minutes.

If the expression returns the integer 0, PingFederate does not store the grant and does not issue a refresh token.

If the expression returns any other value, PingFederate sets the lifetime of the persistent grant based on the per-client Persistent Grants Max Lifetime setting.

To set a static lifetime, select **Text** from the **Source** list and enter a static value in the **Value**

This is suitable for testing purposes, or cases where the persistent grant lifetime must always be set to a specific value.

Expression

This option provides more complex mapping capabilities, such as transforming incoming values into different formats. All of the variables available for text entries are also available for expressions.



Tip:

If you need to map two attribute values from a SCIM request to one LDAP attribute value, use an OGNL expression to create the LDAP attribute.



Enable OGNL expression by editing the <pf install>/pingfederate/server/default/ data/config-store/org.sourceid.common.ExpressionManager.xml file. Restart PingFederate after saving the change.

For a clustered PingFederate environment, edit the

org.sourceid.common.ExpressionManager.xml file on the console node, sign on to the administrative console to replicate this change to all engine nodes in the System # Server # Cluster Management window, and restart all nodes.

SCIM Group

When you make this selection, the associated Value list populates with the defined components of the SCIM request.

No Mapping

Select this option to ignore the Value field.

Text

The value is what you enter. This can be text only, or you can mix text with references to any of the values from the SCIM request, using the \${attribute} syntax.

2. Click Done.

Reviewing group mapping (Write Groups) configuration

You can review your group mapping settings before completing configuration.

About this task

The **Summary** tab provides an overview of the inbound provisioning configuration for request mapping.

Steps

- To amend your configuration, click the corresponding tab title, then follow the configuration wizard to complete the task.
- To keep your changes, click **Done** and continue with the rest of the configuration.



When editing an existing configuration, you can also click **Save** as soon as the administrative console offers the opportunity to do so.

• To discard your changes, click **Cancel**.

Configuring a SCIM response for groups

To configure a System for Cross-domain Identity Management (SCIM) response to a request to read and return provisioned SCIM attributes, identify and map the group attributes you want to include.

About this task

You can begin to identify and map these attributes on the Read Groups tab.



Steps

Click Configure Read Groups to continue.

Identifying expected group attributes for the SCIM response

On the Attribute Contract tab, you can identify the group attributes for you and your partner.

About this task

An attribute contract is a set of group attributes that you and your partner have agreed will be sent in a System for Cross-domain Identity Management (SCIM) response for this connection. The attributes you mapped to group attributes on the **Write Groups** tab appear at the top of the tab.



Click **Available SCIM Attributes** near the lower-left corner of the tab to include additional attributes you want to map in the SCIM response.

Optionally, you can mask the values of attributes in the log files that PingFederate writes when it sends the SCIM response.

There are several SCIM attributes that are managed internally by PingFederate and are unavailable for inclusion in the attribute contract:

- id
- members

Steps

To add an attribute, enter the attribute name in the text box, select the check box under Mask Values in Log as needed, and click Add.



Attribute names are case-sensitive and must correspond to the attribute names expected by your partner. To see a list of available attributes, click **Available SCIM attributes**.

Choice	Action
Modify an attribute name or masking selection	Click Edit under Action for the attribute, make the change, then click Update .
	Note:
	If you change your mind, ensure that you click Cancel under Actions and not the Cancel button, which discards any other changes you might have made in the configuration steps.
Delete an attribute	Click Delete under Action for the attribute.

Identifying LDAP group attributes for the SCIM response

On the Attributes tab, select the LDAP attributes you want to map to attributes in the System for Crossdomain Identity Management (SCIM) response.

About this task





Note:

This tab only appears if you are configuring an LDAP user store for provisioning and the User and Group Support option is selected on the Connection Type tab.

Steps

- 1. Select a root object class and an attribute from the lists, and then click Add Attribute.
- 2. Repeat the previous step for each attribute requiring provisioning.

Map outgoing group attributes to System for Cross-domain Identity Management (SCIM) responses to READ requests.

About this task



Steps

- 1. On the Attribute Fulfillment tab, for each attribute, select a source from the Source list and then choose or enter a value. You must map all attributes.
 - Context

When selected, the Value list populates with the available context of the transaction. Select the desired context from the list.



Note:

As the HTTP Request context value is retrieved as a Java object rather than text, use OGNL expressions to evaluate and return values.



If you are configuring an OAuth Attribute Mapping configuration and have added PERSISTENT GRANT LIFETIME as an extended attribute in the Authorization Server Settings window, you can set the lifetime of persistent grants based on the outcome of attribute mapping expressions or the per-client Persistent Grants Max Lifetime Setting.

- To set lifetime based on the per-client Persistent Grants Max Lifetime setting, select Context from the Source list and Default Persistent Grant Lifetime from the Value list.
- To set lifetime based on the outcome of attribute mapping expressions, select Expression as the source and enter an OGNL expression in the Value field.

If the expression returns a positive integer, the value represents the lifetime of the persistent grant in minutes.

If the expression returns the integer 0, PingFederate does not store the grant and does not issue a refresh token.

If the expression returns any other value, PingFederate sets the lifetime of the persistent grant based on the per-client Persistent Grants Max Lifetime setting.

To set a static lifetime, select Text from the Source list and enter a static value in the Value

This is suitable for testing purposes, or cases where the persistent grant lifetime must always be set to a specific value.

Expression



Tip:

Enable OGNL expression by editing the cpf install/pingfederate/server/default/ data/config-store/org.sourceid.common.ExpressionManager.xml file. Restart PingFederate after saving the change.

For a clustered PingFederate environment, edit the

org.sourceid.common.ExpressionManager.xml file on the console node, sign on to the administrative console to replicate this change to all engine nodes in the System # Server # Cluster Management window, and restart all nodes.

This option provides more complex mapping capabilities, transforming outgoing values into different formats. All of the variables available for text entries are also available for expressions.



If an LDAP attribute needs to be mapped to two attributes in a SCIM response, use an OGNL expression to create them.

LDAP

Values are returned from your query. When you make this selection, the **Value** list populates with the LDAP attributes you identified for this datastore.

Identity Store

Values are returned from your query. When you make this selection, the Value list populates with the Identity Store attributes you identified for this datastore.

No Mapping

Select this option to ignore the **Value** field.

Text

The value is what you enter. This can be text only, or you can mix text with references to any of the values from the SCIM request, using the \${attribute} syntax.

2. Click Done.

Reviewing SCIM response for groups (Read Groups) configuration

You can review your System for Cross-domain Identity Management (SCIM) settings before completing configuration.

About this task

The Summary tab provides an overview of the inbound provisioning configuration for SCIM response mapping.

Steps

- To amend your configuration, click the corresponding tab title, then follow the configuration wizard to complete the task.
- To keep your changes, click **Done** and continue with the rest of the configuration.



Z Tip:

When editing an existing configuration, you can also click Save as soon as the administrative console offers the opportunity to do so.

To discard your changes, click Cancel.

Reviewing the inbound provisioning configuration

You can review your inbound provisioning settings before completing the configuration.

Steps

- To amend your configuration, click the corresponding tab title, then follow the configuration wizard to complete the task.
- To keep your changes, click **Done** and continue with the rest of the configuration.



When editing an existing configuration, you can also click **Save** as soon as the administrative console offers the opportunity to do so.

To discard your changes, click Cancel.

Configuring security credentials

You can configure security credentials and requirements as needed.

About this task

The **Credentials** tab provides the launching point for configuring security requirements you might need depending on the federation protocol you are using and the choices you have made.

Steps

To continue, click Configure Credentials.

For more information and configuration steps see:

- Configuring back-channel authentication for outbound messages on page 732
- Configuring back-channel authentication for inbound messages on page 733
- Managing digital signature settings on page 735
- Managing signature verification settings on page 735
- Choosing an encryption certificate (SAML 2.0) on page 737
- Choosing a decryption key (SAML 2.0) on page 738
- Reviewing IdP credential settings on page 738
- Reviewing an IdP connection on page 739

IdP connection management

When configuring a profile for the inbound artifact binding, outbound SOAP binding, or provisioning, you must specify back-channel authentication information for sending SOAP messages, artifact resolution requests, and provisioning requests to your partner identity provider (IdP).

Similarly, if you send artifacts, SOAP messages, or provisioning messages to your partner IdP, then you must configure SOAP authentication requirements for receiving SOAP responses, artifact resolution requests, or provisioning requests from your partner.

Back-channel authentication also applies to attribute-request configurations, because this profile always uses the SOAP back channel.

See Configuring back-channel authentication for outbound messages on page 732 and Configuring backchannel authentication for inbound messages on page 733 for configuration steps.

Configuring back-channel authentication for outbound messages

You can add and edit configuration settings for back-channel authentication for outbound messages.

Steps

- 1. On the Back-Channel Authentication tab, in the Send to your partner section, click Configure.
- 2. On the **Outbound SOAP Authentication Type** tab, choose one or more authentication methods.

HTTP Basic

When selected, the administrative console prompts you to enter the credentials on the Basic SOAP Authentication (Outbound) tab.

You must obtain these credentials from your partner.

SSL Client Certificate

Applicable only if you specify an endpoint that uses HTTPS.

When selected, the administrative console prompts you to specify your client certificate on the SSL Authentication Certificate tab. If you have not yet created or imported the client certificate, click Manage Certificates to do so. For more information, see .



Important:

When exporting this client certificate for your partner, choose the Certificate Only option.

Digital Signature (Browser SSO profile only)

You select a signing certificate on the **Digital Signature Settings** tab.

This option leverages on the digital signature of the message.

Perform validation on partner's SSL server certificate when SSL used

By default, PingFederate validates your partner's HTTPS server certificate, verifying that the certificate chain is rooted by a trusted certificate authority (CA) and that the hostname matches the certificate's common name (CN).

Clear the associated check box if you do not want this validation to occur.

These options can be used in any combination or independently.

- 3. On the **Summary** tab, review your configuration and perform one of the following tasks: Choose from:
 - Amend your configuration by clicking the corresponding tab title, then follow the configuration wizard to complete the task.
 - Keep your changes by clicking **Done** and continue with the rest of the configuration.



Tip:

When editing an existing configuration, you can also click **Save** as soon as the administrative console offers the opportunity to do so.

Discard your changes by clicking Cancel.

Configuring back-channel authentication for inbound messages

You can add and edit configuration settings for back-channel authentication for inbound messages.

Steps

- 1. On the Back-Channel Authentication tab, in the Received from your partner section, click Configure.
- 2. On the **Inbound Authentication Type** tab, choose one or more authentication methods.

HTTP Basic

When selected, the administrative console prompts you to enter the credentials on the **Basic** SOAP Authentication (Inbound) tab.



Important:

If you are configuring more than one connection that uses the artifact or HTTP profile, you must ensure that the username is unique for each connection. You must communicate these credentials to your partner out-of-band.

SSL Client Certificate

When selected, the administrative console prompts you to specify the trust model and the related certificate settings on subsequent windows. See the next step.

Digital Signature (Browser SSO profile only)

You select a signing certificate on the **Signature Verification Settings** tab.

This option leverages on the digital signature of the message.

Require SSL

When selected, incoming HTTP transmissions must use a secure channel. This option is selected by default.

For SAML 2.0, use these options in any combination or independently. For SAML 1.x, you must enable HTTP Basic authentication, client certificate authentication, or both. You can also add digital signing to ensure message integrity.

3. If you chose SSL Client Certificate in the previous step, select a trust model on the Certificate Verification Method tab.

Anchored

The partner certificate must be signed by a trusted certificate authority (CA). Optionally, you can also restrict the issuer to a specific Trusted CA to mitigate potential man-in-the-middle attacks and to provide a means to isolate certificates used by different connections. The CA's certificate must be imported into the PingFederate Trusted CA store on the Trusted CAswindow...

Unanchored

The partner certificate is self-signed or you want to trust a specified certificate.



Note:

When anchored certificates are used between partners, certificates can be changed without sending the update to your partner. If the certificate is unanchored, any changes must be promulgated.

For more information, see .

Trust model	Subsequent steps
Anchored	On the Subject DN tab:
	 a. Enter the Subject DN of the certificate. b. Optionally, select the Restrict Issuer check box and enter the Issuer DN of the certificate.
	Important: Consider enabling this option to mitigate potential man-in-the-middle attacks and to provide a means to isolate certificates used by different connections.
Unanchored	On the SSL Verification Certificate tab, select the client certification from your partner.
	If you have not yet imported the client certificate from your partner, click Manage Certificates to do so. For more information, see .

- 4. On the **Summary** tab, review your configuration and perform one of the following tasks: Choose from:
 - Amend your configuration by clicking the corresponding tab title, then follow the configuration wizard to complete the task.
 - Keep your changes by clicking **Done** and continue with the rest of the configuration.



When editing an existing configuration, you can also click Save as soon as the administrative console offers the opportunity to do so.

• Discard your changes by clicking **Cancel**.

Managing digital signature settings

Managing digital signature settings defines the private key you will use to sign single sign-on (SSO) authentication or attribute requests (optionally) or SAML 2.0 single logout (SLO) messages for this identity provider (IdP).

About this task

This process allows you to include Key Info with the XML message if you and your partner have agreed to this option.

Digital signing applies to service provider (SP)-initiated SSO under SAML 2.0, when specified by your partner agreement, and to either SLO profile using the POST or redirect bindings. Digital signing also applies if you are configuring an Attribute Query profile and have specified that you will sign attribute requests.

The step is not required for SAML 1.x IdP connections.

Steps

- 1. On the **Digital Signature Settings** tab, select a signing certificate from the **Signing Certificate**list.
 - If you have not yet created or imported your certificate into PingFederate, click Manage Certificates. For more information, see *Manage digital signing certificates and decryption keys* on page 608.
- 2. Optional: Select the Include the certificate in the signature <KeyInfo> element check box if you have agreed to send your public key with the message.
 - Select the Include the raw key in the signature <KeyValue> element check box if your partner agreement requires it.
- 3. Optional: Select the signing algorithm from the list.
 - The default selection is RSA SHA256 or ECDSA SHA256, depending on the Key Algorithm value of the selected digital signing certificate. Make a different selection if you and your partner have agreed to use a stronger algorithm.

Managing signature verification settings

Under SAML 2.0 specifications, when your site receives any SAML 2.0 messages with the POST or Redirect bindings, the messages must be digitally signed.

About this task

Signing is also always required for the SAML 1.x POST binding and for WS-Federation assertions, as well as incoming SAML 1.1 or 2.0 tokens for WS-Trust STS processing.

Depending on your agreement with this idenity provider (IdP), single sign-on (SSO) assertions, SAML 2.0 artifacts, or SOAP messages might also require signatures.

- 1. On the Signature Verification Settings tab, click Manage Signature Verification Settings.
- 2. On the Trust Model tab, select a trust model on the Certificate Verification Method tab.

Anchored

The partner certificate must be signed by a trusted certificate authority (CA). Optionally, you can also restrict the issuer to a specific Trusted CA to mitigate potential man-in-the-middle attacks and to provide a means to isolate certificates used by different connections. The CA's certificate must be imported into the PingFederate Trusted CA store in the Security # Certificate & Key Management # Trusted CAs window.



Important:

If you are using the redirect binding for single logout (SLO) or establishing an OAuth assertion grant connection to exchange JSON web tokens (JWTs) for access tokens. you cannot use anchored certificates because SAML 2.0 does not permit certificates to be included using this transport method and the signature verification process for JWTs requires the public keys to validate the digital signatures.

Unanchored

The partner certificate is self-signed or you want to trust a specified certificate.



Note:

When anchored certificates are used between partners, certificates can be changed without sending the update to your partner. If the certificate is unanchored, any changes must be promulgated.

For more information, see *Digital signing policy coordination* on page 88.

Trust model	Subsequent steps
Anchored	On the Subject DN window:
	 a. Enter the Subject DN of the certificate or extract it from your service provider (SP) partner's certificate if the certificate is stored on an accessible file system. b. (Optional) Select the Restrict Issuer check box and enter the Issuer DN of the certificate. Alternatively, extract it from your partner's certificate.
	Important: Consider enabling this option to mitigate potential man-in-the-middle attacks and to provide a means to isolate certificates used by different connections.

Trust model	Subsequent steps
Unanchored	On the Signature Verification Certificate window:
	a. Select a certificate from the list.
	If you have not yet imported the certificate from your partner, click Manage Certificates to do so. See Managing certificates from partners on page 622. b. (Optional) Select additional certificates.
	Note:
	When configured, PingFederate considers a digital signature valid so long as it can verify the signature using one of the certificates from this list.
	Tip:
	This is useful in situations where your partner has sent you a certificate to replace the current certificate. Adding this second certificate allows PingFederate to continue validating digital signatures as the partner switches to the new signing certificate.
	It also adds support for the scenario where your partner uses a pool for certificates to sign its messages. Adding these certificates ensures digital signatures can be validated as the partner rotates its signing certificates.

3. On the **Summary** tab, review your configuration and perform one of the following tasks.

Amend your configuration

Click the corresponding tab title and then follow the configuration wizard to complete the task.

Keep your changes

Click **Done** and continue with the rest of the configuration.



When editing an existing configuration, you can also click Save as soon as the administrative console offers the opportunity to do so.

Discard your changes

Click Cancel.

Choosing an encryption certificate (SAML 2.0)

If SAML SUBJECT is encrypted, either by itself or as part of a whole assertion, then all references to this name identifier in SAML 2.0 single logout (SLO) requests from your site might also be encrypted if the connection uses service provider (SP)-initiated SLO.

About this task

You must also choose a certificate if encryption of the name identifier is required for an Attribute Request profile. For more information, see Specifying XML encryption policy (for SAML 2.0) on page 696.

Steps

1. Optional: Select an option under **Block Encryption Algorithm**.



Important:

Due to the import restrictions of some countries, Oracle Server Java SE Runtime Environment (JRE) 8 has built-in restrictions on available cryptographic strength (key size). To use larger key sizes, enable the Java Cryptography Extension (JCE) unlimited strength jurisdiction policy. For more information, see the Java 8 release notes in Oracle's documentation.

For Oracle Java SE Development Kit 11, the JCE jurisdiction policy defaults to unlimited strength. For more information, see the Oracle JDK Migration Guide in Oracle's documentation.

The default selection is AES-128.

For more information about XML block encryption and key transport algorithms, see XML Encryption Syntax and Processing from W3C.

2. Select an option under Key Transport Algorithm.



Note:

Due to security risks associated with the RSA-v1.5 algorithm used for key transport, it is no longer available for new connections. Existing connections in which this algorithm is configured continue to support it. However, you should upgrade such connections to use the newer algorithm RSA-OAEP.

The default selection is RSA-OAEP.

3. Select a partner certificate from the list.

If you have not imported the certificate from your partner, click Manage Certificates to do so. For more information see .

Choosing a decryption key (SAML 2.0)

As part of XML encryption, you must identify a certificate and key for PingFederate to use to decrypt incoming assertions or assertion elements.

About this task

For more information on XML encryption, see Specifying XML encryption policy (for SAML 2.0) on page 696.

Steps

1. Select the primary XML decryption key from the list.

If you have not created or imported your certificate into PingFederate, click Manage Certificates. For more information, see .

2. Optional: Select the secondary XML decryption key from the list.

Reviewing IdP credential settings

You can review chosen credential settings before completing configuration.

Steps

 To amend your configuration, click the corresponding tab title, then follow the configuration wizard to complete the task.

To keep your changes, click Done and continue with the rest of the configuration.



When editing an existing configuration, you can also click Save as soon as the administrative console offers the opportunity to do so.

• To discard your changes, click Cancel.

Reviewing an IdP connection

When you finish creating or modifying a connection, you can review the connection settings and toggle the connection status on the Activation & Summary tab.

About this task



Important:

When creating a new connection, the default connection status is **Enabled** when you reach the **Activation** & Summary tab.

Regardless of whether you choose to disable a new connection now or later, you must click Save on the **Activation & Summary** tab if you want to keep the new connection.

Steps

- To amend your configuration, click the corresponding tab title and then follow the steps to complete the task.
- To keep your changes, click Save.
- To discard your changes, click Cancel.

Result

The SSO Application Endpoint provides a sample URL at the /sp/startSSO.ping application endpoint that webmasters or web application developers at your site might use to invoke single sign-on (SSO) for the connection. For a list of supported parameters, see Viewing SP application endpoints on page 657.

If you have selected the **No Mapping** option on the **Identity Mapping** tab, the **Summary & Activation** tab does not show the SSO Application Endpoint sample URL.

OpenID Connect Relying Party support

PingFederate can leverage identities from OpenID Providers (OPs) to complete browser single sign-on (SSO) requests.

In this use case, PingFederate is an OAuth client, specifically a relying party (RP) to the OP. PingFederate supports both the Basic Client and the Implicit Client profiles.

The setup involves establishing an IdP connection to the OP. PingFederate retrieves identity information from the OP and passes the end-user claims, which are user attributes in an ID token, to one or more target applications. This configuration allows administrators to take advantage of their existing last-mile integration and expand the horizon of their applications to additional partners using the OpenID Connect protocol.

If the OP supports the OpenID Connect Discovery specification, the connection setup is expedited by loading the metadata from the OP. Based on the discovery specification, PingFederate makes a direct HTTP GET request to the /.well-known/openid-configuration endpoint at the OP and populates the attribute contract and the protocol settings of the connection automatically. Manual adjustments can be made during the connection setup or at a later time.

Additionally, you can refresh the protocol settings by reloading the metadata from the OP at any time. If additional claims are supported, PingFederate adds them to the attribute contract, so that they can be mapped to the target applications later. In the event that the previously supported claims have been dropped by the OP from the metadata, they remain in the attribute contract. While the existing mapping configuration might not be adversely affected, runtime errors might occur if certain attributes are no longer available to the target applications. If runtime errors occur, review the server log and modify the mapping configuration accordingly.

If applicable, administrators can define additional request parameters, which can be included in the authentication requests to support OP-specific use cases. Administrators can restrict the values to those defined in the configuration. Alternatively, administrators can allow the target applications to optionally override the values at runtime. As an added security measure, administrators can protect the requested authentication context, acr values, and authentication requirement, prompt, so that these parameters in the authentication requests cannot be overridden by the target applications. By default, PingFederate sends these request parameters through multiple query parameters unsigned. Optionally, administrators can configure PingFederate to send request parameters through a request object by value, in which case request parameters are represented by individual claims in a signed JSON Web Token (JWT). When the OP receives the authentication request, it can validate the integrity of the request parameters based on the digital signature in the JWT. For more information, see the section explaining passing a request object by value in the OpenID Connect specification at openid.net/specs/openid-connectcore-1 0.html#RequestObject.

Processing steps

- 1. A user starts a browser SSO request at the /sp/startSSO.ping service provider (SP) application endpoint or the /sp/init login.ping SP protocol endpoint.
- 2. The relying party, PingFederate, sends to the OP an authentication request though the browser containing the desired request parameters, such as response type, scope and redirect uri, with or without any custom query parameters.
- 3. The OP prompts the user to authenticate and authorize, as needed.
- 4. The OP sends the user back to PingFederate, at the redirect uri parameter value, through the browser with an authorization code for the Basic Client profile, or an ID token and an access token if specified in the configuration for the Implicit Client profile.
- 5. Applicable only to the Basic Client profile, PingFederate sends a token request with the authorization code to the OP at its token endpoint, directly through a back-channel HTTP POST request. The OP returns to PingFederate an access token and an ID token.
- 6. PingFederate validates the ID token.
- 7. PingFederate passes the end-user claims to the target application through an SP adapter instance or an authentication policy contract through the browser.

The access token, if any, can also be passed to the target application, so that API security use cases can be layered on top of the browser-based SSO request.



If the UserInfo endpoint is included as part of the connection settings and an access token is provided by the OP, PingFederate also retrieves claims from the OP before the last step.

For more information about OpenID Connect, see openid.net/connect.

You can create an OpenID Connect identity provider (IdP) connection to take advantage of your existing last-mile integration and expand the horizon of your applications to additional partners using the OpenID Connect protocol.

Steps

- 1. Go to Authentication # Integration # IdP Connections, and then create a new IdP connection.
- 2. On the Connection Type tab, select the Browser SSO Profiles check box, and from the Protocol list, select OpenID Connect.



Note:

When OpenID Connect is the chosen protocol, the other types become unavailable.

3. On the Connection Options tab, you can enable just-in-time (JIT) provisioning, OAuth attribute mapping, which requires the OAuth 2.0 authorization server role, or both.



Note:

For simplicity, this topic focuses on managing OpenID Connect IdP connection settings.

- 4. On the General Info tab:
 - a. Provide the required information, which includes:

Issuer

The Issuer Identifier of the OpenID Provider (OP).

Connection Name

A plain-language identifier for the connection; for example, a company or department name. This name is displayed in the connection list on the administrative console.

Client ID

The client ID to communicate with the OP.

This client represents PingFederate and is created and managed at the OP. For more information, see the documentation provided by the OP.

Client Secret

The client secret to communicate with the OP.

Applicable only when the client representing PingFederate supports the Basic Client profile. For more information, see step 12.

b. Optional: Click Load Metadata.



Loading metadata from the OpenID Provider (OP) expedites the connection setup. You can also update an existing connection by reloading metadata.

- 5. On the Browser SSO tab, click Configure Browser SSO.
- 6. On the User-Session Creation tab, click Configure User-Session Creation.

- 7. On the **Identity Mapping** tab, you have three choices: Choose from:
 - Select the No Mapping check box if you plan on passing end-user claims to the target application through an authentication policy contract in an SP authentication policy.
 - Select the Account Mapping check box if you plan on passing end-user claims to the target application through an SP adapter instance or an authentication policy contract if your PingFederate server is a federation hub that bridges an OP to an SP.
 - Select the Account Linking check box if your target application requires account linking.



End-user claims are basically user attributes found in ID tokens or obtained from the UserInfo endpoint at the OP.

For illustration, this topic uses the **Account Mapping** configuration.

8. On the **Attribute Contract** tab, extend the attribute contract.

To mask the attribute values in the log, select the relevant check box for each applicable end-user claim.



Note:

If you have chosen to load the metadata from the OP on the General Info tab, the attribute contract is populated automatically.

9. On the Target Session Mapping tab, click Map New Adapter Instance to map end-user claims to the target application through an SP adapter instance or an authentication policy contract.

Follow the administrative console to fulfill the SP adapter contract or the authentication policy contract. Like other IdP connections, you can query additional attributes from a datastore, specify issuance criteria, or both. When mapping an attribute, select Provider Claims from the Source list to map the attribute to an end-user claim.

If your target application requires the associated access token, select Context as the source and Access Token as the value.



Note:

If the client representing PingFederate supports the Basic Client profile, PingFederate always receives an access token from the OP to retrieve an ID token.

If the client supports the Implicit Client profile, you must select the Form POST with access token option in step 12, such that the OP will return an access token and an ID token as part of the authentication and authorization flow.

The Target Session Mapping configuration does not apply when the No Mapping option is selected on the Identity Mapping tab.

10. On the Protocol Settings tab, click Configure Protocol Settings.

11. On the **OpenID Provider Info** tab, provide the scopes, the endpoints, and the authentication scheme. See the following image for an example.





If you have chosen to load the metadata from the OpenID Provider (OP) on the General Info tab, the **Scopes** field and all endpoints are pre-populated, provided that the metadata contains the information.

Field	Description
Scopes	The scopes to be included in the OpenID Connect authentication and OAuth token requests to the OP. Multiple space-separated values are allowed.
	The default value, without loading metadata from the OP, is openid.
	Tip:
	For a list of OpenID Connect defined scopes, see the section about requesting claims using scope values in the OpenID Connection specification at openid.net/specs/openid-connect-core-1_0.html#ScopeClaims.
Authorization Endpoint	The authorization endpoint at the OP.
	You can enter a relative path, starting with a forward slash, if you provide base URL on the General Info tab.
	There is no default value without loading metadata from the OP.
OpenID Connect Login Type	The OpenID Connect client profile of the client. This client represents PingFederate and is created and managed at the OP.
	 If the client is configured to support the Basic Client profile, select Code.
	The resulting value of the response_type parameter is code. If the client is configured to support the Implicit Client profile, select Form POST.
	The resulting value of the response_type parameter is id_token. If the client is configured to support the Implicit Client profile and the target application requires the associated access token, select Form POST with access token.
	The resulting values of the response_type parameter are id_token token.
	The default selection, without loading metadata from the OP, is Code.

Field

Description

Authentication Scheme

The client authentication method that PingFederate uses. Applicable and visible only to clients supporting the Basic Client profile.

- Select Basic to submit credentials with HTTP Basic authentication.
- Select **POST** to submit credentials with POST.
- Select **Private Key JWT** to authenticate with the private_key_jwt Client Authentication method, see Client Authentication in the OpenID Connect specification.

The default selection, without loading metadata from the OP, is **Basic**.

Authentication Signing Algorithm

Select the algorithm that PingFederate uses to sign the JWT.

Applicable and visible only when **Private Key JWT** is the chosen authentication scheme.

If PingFederate is either deployed to run in a Java 11 runtime environment or integrated with a hardware security module (HSM) and configured to use static keys for OAuth and OpenID Connect, additional RSASSA-PSS signing algorithms become available for selection. For more information on HSM integration and static keys, see and .



Note:

If static keys for OAuth and OpenID Connect are enabled, Elliptic-curve cryptography (EC) algorithms that have not been configured with an active static keys are hidden.

Changes made in the static-key configuration might affect runtime transactions and require additional changes here. For more information, see.



Based on the chosen signing algorithm, PingFederate selects the signing JSON Web Key (JWK) from its JWK Set (JWKS) at runtime.

In order for the OP to validate the signed JWT, ensure that the OP can access your PingFederate JWKS endpoint, which returns the current set of JSON Web Keys. The PingFederate JWKS endpoint is located at <Base URL>/pf/JWKS, where Base URL is defined on System # Server # Protocol Settings # Federation Info.

For example, if the Base URL field value is https://www.example.com, the PingFederate JWKS endpoint is https://www.example.com/pf/JWKS. You can pass the PingFederate JWKS endpoint directly to the OP or have the OP contact the PingFederate OpenID Provider configuration endpoint to obtain the information. For more information, see .

Ciold	Description
Field	Description
Token Endpoint, UserInfo Endpoint,	Various OAuth 2.0 and OpenID Connect 1.0 endpoints at the OP. For more information, see openid.net/connect.
and JWKS URL	Token Endpoint
	The Token Endpoint field is only visible and required for clients supporting the Basic Client profile. (In other words, the OpenID Connect Login Type field is set to Code .)
	UserInfo Endpoint
	The UserInfo Endpoint field is optional. If omitted, PingFederate only has access to the end-user claims from the ID tokens.
	JWKS URL
	The JWKS URL is required in order for PingFederate to validate the inbound ID tokens from the OP. If the OP signs its JWTs using an RSASSA-PSS signing algorithm, PingFederate must be deployed to run in a Java 11 runtime environment or integrated with a hardware security module (HSM) to process the digital signatures. For more information on HSM integration, see .
	There are no default values without loading metadata from the OP.
Sign Request	Select this check box to send request parameters as claims in a request object, a self-contained, signed JWT as one request query parameter to the OP.
	When this optional configuration is enabled, the OP can validate the integrity of the request parameters based on the digital signature found in the signed JWT. For more information, see the section explaining passing a request object by value in the OpenID Connect specification at openid.net/specs/openid-connect-core-1_0.html#RequestObject.
	This check box is not selected by default, in which case PingFederate sends request parameters with multiple query parameters, unsigned.

Field

Description

Request Signing Algorithm

Select the algorithm that PingFederate uses to sign the request object.

Applicable and visible only when the **Sign Request** check box is selected.

If PingFederate is either deployed to run in a Java 11 runtime environment or integrated with a hardware security module (HSM) and configured to use static keys for OAuth and OpenID Connect, additional RSASSA-PSS signing algorithms become available for selection. For more information on HSM integration and static keys, see and .



Note:

If static keys for OAuth and OpenID Connect are enabled, Elliptic-curve cryptography (EC) algorithms that have not been configured with an active static keys are hidden.

Changes made in the static-key configuration might affect runtime transactions and require additional changes here. For more information, see.



Note:

PingFederate automatically selects the signing JSON web key (JWK) based on the selected signing algorithm from its JWK Set (JWKS).

In order for the OP to validate the signed request object, ensure that the OP can access your PingFederate's JWKS URL, which returns the current set of JSON web keys. The PingFederate JWKS URL is located at <Base URL>/pf/JWKS, where Base URL is defined on System # Server # **Protocol Settings # Federation Info.**

For example, if the **Base URL** field value is https://www.example.com, the PingFederate JWKS URL is https://www.example.com/pf/JWKS. You can pass the JWKS URL directly to the OP or have the OP contact the PingFederate OpenID Provider configuration endpoint for it. For more information, see .

- 12. Optional: Remain on the OpenID Provider Info tab and specify the request parameters that are allowed to be included in the authentication requests to the OP under Request Parameters. For more information, see .
- 13. Optional: On the **Overrides** tab, specify a default target URL and authentication context overrides.
- 14. On the **Activation & Summary** tab, review your connection settings.

When you finish setting up a connection, you can choose to activate it immediately.



Important:

Regardless of whether you choose to activate a new connection now or later, you must click Save on the **Summary & Activation** tab for a new connection if you want to keep the configuration.

You can deactivate a connection at any time. When a connection is inactive, all transactions to or from this partner are disabled.

In this use case, because PingFederate is an OAuth client, you are likely required by the authorization server at the OP to register the Redirect URI as shown on the Summary & Activation tab. This

registration should be associated with the client that represents PingFederate, the client that you have provided on the **General Info** tab. For more information, see the documentation provided by the OP.

The **SSO Application Endpoint** provides a sample URL at the /sp/startSSO.ping application endpoint that webmasters or web application developers at your site can use to invoke SSO for the connection. For a list of supported parameters, see *Viewing SP application endpoints* on page 657.

If you have selected the **No Mapping** option on the **Identity Mapping** tab, the **Summary & Activation** tab does not show the **SSO Application Endpoint** sample URL.

The target application can also invoke SSO requests by contacting the <code>/sp/init_login.ping</code> SP protocol endpoint. For more information, see *Configuring request parameters and SSO URLs* on page 747.

Configuring request parameters and SSO URLs

On the **OpenID Provider Info** tab, administrators can define request parameters under **Request Parameters**.

About this task

You can define request parameters for the following purposes:

- Allow custom request parameters to be include in the authentication requests to support OpenID provider (OP)-specific use cases.
- Define the default values for the request parameters.
- Specify whether the default values, if any, can be overridden at runtime.
- Allow the target application to request different scopes at runtime. The OP can reject the requested scopes based on its client configuration.
- Protect the requested authentication request, acr_values, the authentication requirement, prompt, or both so that none of them can be overridden at runtime by the application endpoint parameters RequestedAuthnCtx, IsPassive, and ForceAuthn.

Steps

- 1. Add a request parameter under Name.
- 2. Define a default parameter value under Value.

This is optional if the target application is allowed to override the parameter value at runtime. When no default value is specified, any value provided by the target application is accepted by the <code>/sp/startSSO.ping</code> service provider (SP) application endpoint. If the target application does not provide the parameter in its single sign-on (SSO) URL and no default value is specified, the parameter is not included in the authentication requests.

This is required if the target application is not allowed to override the parameter value at runtime.

When specified, the request parameter is always included in the authentication requests. If the target application is not allowed to override the parameter value at runtime, the default value is sent.

3. Select the check box under **Application Endpoint Override** if the target application is allowed to override the parameter value at runtime.

Result:

If the target application does not provide the parameter in its single sign-on (SSO) URL and the configuration does not include a default value, the parameter is not included in the authentication requests.

If the target application does not provide the parameter in its SSO URL, the default value, if any, is

If the target application provides the parameter in its SSO URL to the /sp/startSSO.ping SP application endpoint, the value in the SSO URL is used.



Note:

The /sp/init login.ping service provider (SP) protocol endpoint does not accept overridden values. The login hint parameter is the only exception. The default value, if any, is used. See the note at the end of this topic for more information.

4. Click Add.

Click Edit, Update, and Cancel to make or undo a change. Click Delete and Undelete to remove an entry or cancel the removal request.

5. Repeat these steps to define another request parameters.

Example

Consider the following sample configuration:



- The hd parameter is defined with a default value that cannot be overridden at runtime. The parameter is always included in the authentication requests and the value is always example.org.
- The customMultiValued parameter is defined with two default values that cannot be overridden at runtime. This multivalued parameter is always included in the authentication requests. The values are always as defined.
- The customOverridableOne parameter is defined with a default value that can be overridden at runtime. This parameter is always included in the authentication requests. If the target application

provides the parameter in its SSO URL, the value in the SSO URL is used. If the target application does not provide the parameter in its SSO URL, the default value is used.

- To override the value, configure the target application to append the request parameter and the desired value to the SSO Application Endpoint, as shown on the Summary & Activation tab, as in the following example.
 - https://sso.example.com/sp/startSSO.ping?PartnerIdpId=https%3A%2F%2Fsso.alpha.local %3A9031 and customOverridableOne=foo
- To construct a multivalued request parameter, append the request parameter multiple times with different values, as in the following example.
 - https://sso.example.com/sp/startSSO.ping?PartnerIdpId=https%3A%2F%2Fsso.alpha.local %3A9031 and customOverridableOne=foo and customOverridableOne=bar
 - https%3A%2F%2Fsso.alpha.local%3A9031 is the URL-encoded value of https:// sso.alpha.local:9031, the issuer value of the OP.
- The customOverridableTwo parameter is defined without a default value. Any value provided by the target application in the SSO URL is accepted. To include this parameter in the authentication requests to the OP, configure the target application to append the request parameter and the desired value to the SSO Application Endpoint.
 - To construct a multivalued request parameter, append the parameter multiple times with different values.
 - If the target application does not provide the parameter in its SSO URL, the parameter is not included in the authentication requests.
- The scope (standard) parameter is defined with a value matching that of the Scopes field, on the same tab, and with the option to allow the target application to override the value at runtime. In essence, the target application is allowed to dynamically change the scope it requires at runtime by appending the scope parameter and the desired scopes to the SSO Application Endpoint.



Note:

While the target application can request different scopes, the OP can reject the requested scopes based on its client configuration. Work with the OP to understand which scopes are applicable to your use case to prevent runtime errors.

- The acr values (standard) parameter is defined with a default value that cannot be overridden at runtime. As a result, the RequestedAuthnCtx parameter, if supplied in the SSO URL by the target application, is ignored. In the authentication requests, the value of the acr values parameter is always set to the default value specified in the configuration. Define the acr values parameter if you want to protect the requested authentication context from the target application.
- The prompt (standard) parameter is defined with a default value of login that cannot be overridden at runtime. As a result, the target application will not be able to suppress the reauthentication requirement by including IsPassive=true in the SSO URL. In the authentication requests, the value of the prompt parameter is always set to login.

Similarly, if the prompt parameter is defined with a default value of none that cannot be overridden at runtime, the target application will not be able to request the end users to reauthenticate by including ForceAuthn=true in the SSO URL. In the authentication requests, the value of the prompt parameter is always set to none.



These examples use the /sp/startSSO.ping SP application endpoint. As needed, you can also use the /sp/init login.ping SP protocol endpoint to invoke the Third Party Initiated Login flow. For more information, see View SP protocol endpoints.



Important:

For information about URL encoding, see third party resources such as HTML URL-encoding Reference.

Query parameters versus request object

By default, PingFederate sends all request parameters through multiple query parameters, unsigned.

If the Sign Request check box is selected, PingFederate creates a signed JSON web token (JWT) that contains claims representing the request parameters and passes the signed JWT as one query parameter, request, to the OpenID provider (OP). The client id, response type, and scope request parameters are always passed to the OP as individual query parameter as well.

Consider the following authentication requests based on the previous sample configuration. The client authenticates through the HTTP Basic authentication scheme and initiates single sign-on (SSO) request without providing overrides for any request parameters.

Request parameters via query parameters

```
https://sso.alpha.local:9031/as/authorization.oauth2
?acr values=PasswordProtectedTransport
&customMultiValued=value+one
&customMultiValued=value+two
&customOverridableOne=value+can+be+overridden
&hd=example.org
&prompt=login
&nonce=ykulMjpwAFk79R1rBOBWm5
&redirect uri=https://www.example.com/sp/
eyJpc3MiOiJodHRwczpcL1wvc3NvLmFscGhhLmxvY2FsOjkwMzEifQ/cb.openid
&state=e75nIlVU6Wa5TMmOwegDPSEI2iO9zd
&client id=RP
&response type=code
&scope=address+phone+edit+openid+profile+admin+email
```

Request parameters via a request object by value

```
https://sso.alpha.local:9031/as/authorization.oauth2
?request=eyJhbG...ZTMifQ.eyJhdW...lJQIn0.IAOpuf...IqCftg
&client id=RP
&response type=code
&scope=address+phone+edit+openid+profile+admin+email
```



Note:

The client id, response type, and scope request parameters are always passed to the OP as individual query parameters as defined in the OpenID Connect specification.

The value of the request query parameter, truncated for readability, is the request object, a signed JWT that contains the request parameters as individual claims, illustrated in the following decoded payload.

```
"aud": "https://sso.alpha.local:9031",
"exp": 1495645410,
"acr values": "PasswordProtectedTransport",
"customMultiValued": [
  "value one",
  "value two"
```

```
"customOverridableOne": "value can be overridden",
   "hd": "example.org",
   "prompt": "login",
   "nonce": "vhW2VJc7eZ6r6vfpiAwepd",
   "redirect_uri": "https://sso.rp.local:9021/sp/
eyJpc3MiOiJodHRwczpcL1wvc3NvLmFscGhhLmxvY2FsOjkwMzEifQ/cb.openid",
   "state": "nFVzgFirZtg3kBXMFpWt5RNhO4oDuA",
   "client_id": "RP",
   "response_type": "code",
   "scope": "address phone edit openid profile admin email"
}
```

For more information, see the section explaining passing a request object by value in the OpenID Connect specification at *openid.net/specs/openid-connect-core-1_0.html#RequestObject*.

Configuring IdP discovery using a persistent cookie

PingFederate's proprietary identity provider (IdP)-discovery method makes use of an IdP persistent reference cookie (IPRC) to track the identity provider with whom a user last authenticated.

About this task

There are three significant differences between standard IdP discovery and the IPRC method:

- Standard IdP discovery can be used only with SAML 2.0, but the IPRC can be used with any federation protocol.
- The common domain cookie (CDC) can be configured as a temporary, session-based cookie. The IPRC always persists for a configurable period of time.
- The CDC is set by the IdP and is readable by both federation partners. The IPRC is set by the service provider (SP), using information in the SAML assertion, and cannot be accessed by the IdP.

The deployed connection configuration between SP and IdP partners must include SP-initiated single signon (SSO).

Steps

- 1. Edit the org.sourceid.websso.profiles.sp.IdpIdCookieSupport.xml file located in the cpf install>/pingfederate/server/default/data/config-store directory.
- 2. Set the value of EnableIdpIdCookie to true.
- Optional: Modify the remaining elements in the configuration, as described in the following table.

Field	Description
IdpldCookieName	The name of the IPRC set by the SP installation. The default is IdPId. The cookie name cannot contain any of the following characters: $\&$, >, <, ; , a comma, or a space.
IdpldCookieLifeTimeInDalys lifetime for the cookie. The default is 365 days and a maximum of 24855 days. The browser will delete the cookie when the period is expired.	
ShowIdpSelectionList	If set to true, the default, the SP displays a list of IdPs that can be used to initiate the SSO event if the cookie is not set. If set to false, the SP installation generates an error page.



Once an IPRC cookie is set, the only way to change the IdP to whom the SP will send Authentication Requests for the user is to do one of the following: wait for the cookie to expire, delete the cookie, or perform IdP-initiated SSO using the new IdP.

System administration

This section describes general administrative functions for PingFederate.

Configuring PingFederate properties

The default administrative console and runtime behavior of PingFederate is controlled in part by configuration properties set in the run.properties file, located in the <pf install>/ pingfederate/bin directory.

About this task

The most common properties are documented in the following table. For the rest of the properties, including various cookie-encoding options, see the run.properties file.



The clustering configuration options are also maintained in the run.properties file. For more information, see *Deploying cluster servers* on page 201.

Property	Description
pf.admin.https.port	Defines the port on which the PingFederate administrative console runs. The default value is 9999.
pf.console.bind.address	Defines the IP address over which the PingFederate administrative console communicates. Use for deployments where multiple network interfaces are installed on the machine running PingFederate.
pf.console.title	Defines the browser window or tab title for the administrative console. It makes separate instances easily identifiable.
pf.console.environment	Defines the name of the PingFederate environment that will be displayed in the administrative console. It makes separate environments easily identifiable.
pf.console.show.background.images	Enables or disables the background images on the dashboard of the administrative console. The images are enabled by default.
pf.console.session.timeout	Defines the length of time in minutes until an inactive administrative console times out. The minimum setting is 1 minute, and maximum is 8 hours (480 minutes). Default is 30 minutes.
pf.log.eventdetail	Enables or disables (the default) detailed event logging for actions performed by administrative console users.

Description
Indicates whether more than one administrative user may access the administrative console at one time. Supported values: Single Multiple. The default value is Multiple.
Indicates whether administrators sign on to PingFederate using credentials managed internally by PingFederate or externally by other systems.
Defines the authentication method of the PingFederate administrative API.
When LDAP administrative console authentication is enabled, indicates the name of the file containing configuration properties.
When certificate-based console authentication is enabled, indicates the name of the file containing configuration properties.
When RADIUS-based console authentication is enabled, indicates the name of the file containing configuration properties.
When OIDC administrative-console authentication is enabled, indicates the name of the file containing configuration properties.
Defines the port on which PingFederate listens for unencrypted HTTP traffic at runtime. For security reasons, this port is turned off by default.
CAUTION:
This port should remain disabled in production if your deployment configuration directly exposes the PingFederate server to the Internet.
Defines the port on which PingFederate listens for encrypted HTTPS (SSL/TLS) traffic. The default value is 9031.
Defines a secondary HTTPS port that can be used for mutual SSL/TLS (client X.509 certificate) authentication for both end users and protocol requests (SAML, WS-Trust, and OAuth). Set its value to the desired inbound listening TCP port. A value of -1 disables this feature.
Important:
If you are using client X.509 certificates for either WS-Trust STS authentication or for SAML back-channel authentication, you must use this port, or a similarly configured new listener, with either the WantClientAuth or NeedClientAuth parameter set to true in the jetty-runtime.xml file.
For more information, see the note at the end of this table.

Property	Description
pf.engine.bind.address	Defines the IP address over which the PingFederate server communicates with partner federation gateways. Use for deployments where multiple network interfaces are installed on the machine running PingFederate.
pf.monitor.bind.address	Defines the IP address over which an SNMP agent and Java Management Extensions (JMX) communicate with PingFederate. Use for deployments where multiple network interfaces are installed on the machine running PingFederate.
pf.engine.prefer_ipv4	Defines the protocol to be used by PingFederate. True, the default, enables use of IPv4 only. False enables use of both IPv4 and IPv6.
http.proxyHost and http.proxyPort	Specifies the hostname, or the IP address, and the port number of the forward proxy server that HTTP traffic originating from PingFederate must go through.
https.proxyHost and https.proxyPort	Specifies the hostname, or the IP address, and the port number of the forward proxy server that HTTPS traffic originating from PingFederate must go through.
http.nonProxyHosts	Specifies one or more destinations where PingFederate is not required to proxy its HTTP and HTTPS traffic through the forward proxy server configured by the http[s].proxyHost and http[s].proxyPort properties. This property supports multiple values separated by the pipe character () and the wildcard character (*) for pattern matching. See the example below.
	*.example.com localhost

Property	Description
pf.runtime.context.path	Allows customization of the server path for PingFederate endpoints.
	Note:
	If this property is changed, the path must also be added to the base URL for your PingFederate environment. Base URL is defined on System # Server # Protocol Settings # Federation Info .
	The pf.runtime.context.path property is also compatible with virtual host names. Unlike the base URL configuration, the virtual host names configuration does not require any context path. Virtual host names are defined on System # Server # Virtual Host Names.
	For example, suppose the base URL is https:// www.example.com:9031 and the virtual host names are www.example.org and www.example.info. To configure the pf.runtime.context.path property value as /sso, you must update the base URL to https://www.example.com:9031/ sso but leave the virtual host names as they are. Once configured, you can access the runtime server at the following endpoints:
	Base URL
	https://www.example.com:9031/sso
	Virtual host names
	https://www.example.org:9031/ssohttps://www.example.info:9031/sso
pf.log.dir	Network path to the output location of log files. The default is
	<pf_install>>/pingfederate/log</pf_install>
pf.hsm.mode	Enables or disables (the default) a FIPS-compliance Hardware Security Module (HSM).
pf.hsm.hybrid	Enables or disables the HSM hybrid mode. Applicable only when the pf.hsm.mode property is configured to use an HSM.
	When set to true, keys and certificates can be stored on either the HSM or the local trust store. When set to false, the default setting, keys and certificates are stored on the HSM when applicable.
	The HSM hybrid mode allows an organization to move the storage of keys and certificates from the local trust store to an HSM over time without deploying a new PingFederate installation and mirroring the setup. For more information, see <i>Transitioning to an HSM</i> on page 626.

Property	Description
org.bouncycastle.fips.approved_only	When the pf.hsm.hybrid property is set to true, this property can be set to true or false. In this case, the recommended setting is false.
	If pf.hsm.hybrid is set to false, this property must be set to true.
	In FIPS-approved mode only, the module will provide approved algorithms only. For more information, see https://www.bouncycastle.org/fips-java/ .
	The default setting is true.
pf.provisioner.mode	Enables or disables (the default) outbound provisioning. Also used to enable provisioning failover.
pf.heartbeat.system.monitoring	Enables or disables (the default) the heartbeat endpoint, / pf/heartbeat.ping, to return detailed system monitoring information through a customizable Velocity template file . For more information, see <i>Customizing the heartbeat message</i> on page 835.
	When set to false, the /pf/heartbeat.ping endpoint returns OK.
	When set to true, the /pf/heartbeat.ping endpoint returns all available stats.
org.apache.xml.security.ignoreLineBreaks	Determines whether PingFederate omits line breaks in XML digital signatures. If omitted, this setting defaults to false. Set this property to true for improved interoperability with Microsoft products.



Note:

Additional configuration of the listener ports, including adding new listeners, is available through the <pf install>/pingfederate/etc/jetty-runtime.xml file. For example, options include the WantClientAuth and NeedClientAuth flags, which indicate that a client certificate is either requested or required, respectively, for mutual SSL/TLS. For the pre-configured SSL secondary port, the WantClientAuth parameter is set to true and the NeedClientAuth parameter is set to false by default.

Steps

- 1. Edit the <pf_install>/pingfederate/bin/run.properties file. You should consider creating a backup copy of the file.
- 2. Modify the applicable properties.
- 3. Restart PingFederate.

Result



Important:

Enabling OIDC-based authentication

You can enable OIDC-based authentication to the administrative console by setting a property in the run.properties file, and by configuring other properties in the oidc.properties file.

About this task



Important:

All endpoints must be HTTPS.

Steps

- 1. On your OIDC provider, configure an OAuth client to represent the PingFederate administrative console. Specify the following redirect URI for the client:
 - https://<pf_admin_hostname>:<pf_admin_port>/pingfederate/app?service=finishsso
 - You need the client's credentials in the following steps.
- 2. Edit the <pf install>/pingfederate/bin/run.properties file, and set the pf.console.authentication property to OIDC.
- 3. Edit the <pf install>/pingfederate/bin/oidc.properties file, and modify the applicable properties as described in the following table.

Property	Description
client.id	The client ID to communicate with the OpenID Provider (OP).
	This property is required.
client.secret	The client secret used to communicate with the OpenID Provider.
	The client secret should be in obfuscated format. We recommend that the secret be obfuscated using one of the following utilities in the / bin directory:
	• On Windows: obfuscate.bat
	• On Linux: ./obfuscate.sh
	For example: obfuscate.bat secret
	This property is required when the client authentication is either client_secret_basic or client_secret_post.
client.authn.method	The type of client authentication that is expected by the token endpoint in the OpenID Provider. Supported values include:
	 client_secret_basic – Client credentials using the HTTP Basic authentication scheme.
	 client_secret_post – Client Credentials included in the request body.
	 private_key_jwt – Client authenticates in accordance with JSON Web Token (JWT).
	This property is required.

Property	Description
authorization.endpoint	The authorization endpoint at the OpenID Provider.
	This property is required.
token.endpoint	The token endpoint at the OpenID Provider. PingFederate OIDC login only supports the authorization code flow.
	This property is required.
user.info.endpoint	The endpoint that is accessed when the required claims are not present in the ID tokens.
	This property is optional.
end.session.endpoint	The end session endpoint at the OpenID Provider. When no value is provided, the administrator will be redirected to the default PingFederate logout page.
	This property is optional.
issuer	The issuer identifier of the OpenID Provider. The value provided is matched with the iss claims in the obtained ID token.
	This property is required.
acr.values	The authentication context class reference values that will be used by the OpenID Provider.
	This property is optional.
scopes	The authorization endpoint at the OpenID Provider. The default setting is openid.
	This property is required.
username.attribute.name	The name of the claim that represents the username of the administrator. The default setting is sub.
	This property is required.
role.attribute.name	The name of the claim that is used to determine the role for administrators.
	This property is required.
role.map.admin. <i>n</i>	Used when multiple values need to be mapped to a single PingFederate role. In this case, multiple properties must be created using a numeric, incremental suffice, starting with 1. For example:
	role.map.admin.1= role.map.admin.2= role.map.admin.3=

Property	Description
role.admin=	The administrator role claim value mapping. For example, assume
role.cryptoManager=	that admin_role is a claim and the possible values for it are:
role.userAdmin=	role.admin=admin role.cryptoManager=crypto
role.expressionAdmin=	role.userAdmin=uadmin
	role.expressionAdmin=eadmin
	When the claim admin_role has more than one value, for example, admin_role: ["admin", "crypto", "uadmin"], the user will be granted admin, crypto, and user administrator roles.
	When the claim admin_role has a single value, for example, admin_role: "admin", the user will be granted admin role.
	This property is required.
role.auditor	The auditor role claim value mapping.
	This property is optional.

4. Restart PingFederate.

Configuring size limits

You can configure size limits, which include values in the IdP Session Registry, the SP Session Registry, the Inter-Request State management service, and others.

Steps

- 1. To configure a size limit, edit the <pf_install>/pingfederate/server/default/conf/size-limits.conf file.
- 2. Modify the applicable settings. The setting are described in the following table.

Setting	Description	
The IdP Session Registry stores SLO-related session information for IdP adapters, as well as for IdP and SP connections that a user has interacted with. It also stores sessions for PingFederate's Authentication Sessions feature.		
IdpSessionRegistryMapImpI.max.sessions	This setting controls the maximum number of user sessions (for SLO or Authentication Sessions) kept in memory. When this limit is reached, sessions are removed on a least-recently-used basis. The default setting is 10000.	
IdpSessionRegistryMapImpl.max.individual.session	nshis setting controls the maximum number of IdP adapter or IdP connection sessions per user session. When this limit is reached, sessions are removed on a first-in first-out basis. The default setting is 500.	

Setting	Description
IdpSessionRegistryMapImpI.max.partner.sessions	This setting controls the maximum number of SP connection sessions per IdP adapter or IdP connection session. When this limit is reached, sessions are removed on a first-in first-out basis.
	The default setting is 500.
IdpSessionRegistryMapImpl.expiry.mins	This setting defines the expiry period for user sessions in minutes. If no activity has been seen for a given user session for this period, it will be removed.
	The default setting is 1440
The SP Session Registry stores SLO-related sessi connections that a user has interacted with.	on information for SP adapters and IdP
SpSessionRegistryMapImpl.max.sessions	This setting controls the maximum number of user sessions kept in memory. When this limit is reached, sessions are removed on a least-recently-used basis.
	The default setting is 10000.
SpSessionRegistryMapImpl.max.individual.session	If his setting controls the maximum number of SP adapter sessions per user session. When this limit is reached, sessions are removed on a first-in first-out basis.
	The default setting is 500.
SpSessionRegistryMapImpl.expiry.mins	This setting defines the expiry period for user sessions in minutes. If no activity has been seen for a given user session for this period, it will be removed.
	The default setting is 1440.
The Inter-Request State Management service has lived state information between requests within an adapters (such as the HTML form adapter) to store	SSO or SLO transaction. The 'attr' map is used by
InterReqStateMgmtMapImpI.max.size.state.map	This setting controls the maximum number of user sessions in the state map. When this limit is reached, sessions are removed on a least-recently-used basis.
	The default setting is 10000.
InterReqStateMgmtMapImpl.expiry.mins.state.map	This setting controls the expiry period for user sessions in the state map. If no activity has been seen for a given user session for this period, it will be removed.
	The default setting is 30.

Setting	Description
InterReqStateMgmtMapImpl.max.size.attr.map	This setting controls the maximum number of user sessions in the attribute map. When this limit is reached, sessions are removed on a least-recently-used basis.
	The default setting is 10000.
InterReqStateMgmtMapImpl.expiry.mins.attr.map	This setting controls the expiry period for user sessions in the attribute map. If no activity has been seen for a given user session for this period, it will be removed.
	The default setting is 1440.
InterReqStateMgmtMapImpl.max.session.attrs	This setting controls the maximum number of attributes stored in the attribute map for a given user session.
	The default setting is 500.
SessionRevocationServiceMapImpl.max.revoked.s	rishis setting controls the maximum number of revoked session identifiers kept in memory. When this limit is reached, revoked identifiers are removed on a first-in first-out basis.
	The default setting is 50000.
MetadataDirectory.max.size.idp.conn.map	This setting controls the maximum number of IdP connections kept in memory. When this limit is reached, connections are removed on a least-recently-used basis.
	The default setting is 10000.
MetadataDirectory.max.size.sp.conn.map	This setting controls the maximum number of SP connections kept in memory. When this limit is reached, connections are removed on a least-recently-used basis.
	The default setting is 10000.
ClientManagerXmlFileImpl.max.size.clients.map	This setting controls the maximum number of OAuth clients kept in memory. When this limit is reached, clients are removed on a least-recently-used basis.
	The default setting is 10000.

3. Restart PingFederate.

PingFederate log files

PingFederate records document server events depending on your configuration preferences.

PingFederate generates these logs that document server events:

admin.log

Records actions performed by administrative console users.

admin-event-detail.log

Records detailed information about each applicable administrative console event performed by administrative console users if detailed event logging is enabled.

admin-api.log

Records actions performed by administrative-API users.

runtime-api.log

Records actions performed by API users using the OAuth Client Management Service, the OAuth Access Grant Management Service, and the Session Revocation API.

transaction.log

Records individual identity-federation runtime transactions at specified levels of detail.

audit.log

Records a selected, configurable subset of transaction log information plus additional details, intended for security-audit and regulatory compliance purposes.

provisioner-audit.log

Records outbound provisioning events, intended for security-audit purposes.

provisioner.log

Records only provisioning activity.

server.log

Records PingFederate runtime and administrative server activities.

init.log

Records only Jetty messages generated prior to PingFederate start up.

These log files are written to the PingFederate log directory. The default location is the <pf install>/ pingfederate/log directory. As needed, administrators can change the log directory by modifying the **pf.log.dir** property in the <pf install>/pingfederate/bin/run.properties file.

Log4j 2 logging service and configuration

PingFederate uses the Log4j 2 logging service to generate its log files.

Configurations are maintained in the log4j2.xml file, located in the <pf install>/pingfederate/ server/default/conf directory.



Note:

The log4j2.xml configuration file is individually managed per PingFederate server. This flexibility allows multiple PingFederate nodes to write different level of messages to different destinations.

If you want all PingFederate servers to use the same logging configuration, manually synchronize the log4j2.xml file across multiple PingFederate servers.

Log levels and verbosity

Log messages are categorized into six log levels:

- 1. FATAL
- 2. ERROR
- 3. WARN
- 4. INFO
- 5. DEBUG
- 6. TRACE

PingFederate only records messages tagged with log level INFO, WARN, ERROR, and FATAL to the server log and the provisioner log. Messages with DEBUG, or TRACE tags, are not recorded to optimize performance. Console logging is also disabled for the same reason.

For troubleshooting purposes, you can adjust the log level to DEBUG in the log4j2.xml file and reenable console logging.



Important:

When you no longer require debug messages and console logging, turn them off.

For the audit log, the provisioner audit log, and the transaction log, any setting lower than INFO (WARN, ERROR, or FATAL) turns logging off.

For more information, see Enabling debug messages and console logging.

Changes, such as adding a Logger or adjusting log levels, are activated within half a minute. You do not need to restart PingFederate.

Fields (and attributes)

You can customize some logs, such as the audit log and the administrative API log, to log additional or less information by modifying their pattern elements. The log4j2.xml file documents available fields inline.



You can configure PingFederate to log user attributes, if they are present, in the audit log, transaction log, and server log. When you require privacy for sensitive user attributes, select the corresponding check boxes under Mask Log Values to mask their values in these logs.

In addition, messages in the audit log and the server log are recorded with a tracking ID, which can be used to identify subsequent, related transactions. The tracking ID can be used for troubleshooting and support purposes, to aggregate and analyze log entries tied to the same original request. The tracking ID, %X {trackingid}, can also be added to the configuration for the transaction log, or removed from the audit log and the server log by modifying the pattern element for the logs in the log4j2.xml configuration file.

Log formats

The audit log and the provisioner audit log can be written in Common Event Format (CEF). Furthermore, the audit log can also be written in a format used in conjunction with Splunk and the Splunk App for PingFederate. The log4j2.xml file comes preset with configuration samples to ease the setup.

Log destinations

The audit log, the provisioner audit log, the provisioner log, and the server log can be written to databases. PingFederate installation includes setup scripts for various tables, located in the <pf install>/pingfederate/server/default/conf/log4j/sql-scripts directory, and configuration samples in the log4j2.xml file.

Log rotation

Most PingFederate-generated log files roll over at midnight each day. The system keeps all of the resulting historical log files. Some log files, such as the audit.log file, the audit-eventdetail.log file (if enabled), the provisioner-audit.log file (when applicable), and the transaction.log, can become quite large, depending on your production load and settings. You might want to back up or remove older files on a routine basis.

The server.log file is rolled over when it reaches 10 MB. Five old log files are kept before the oldest file is removed. Administrators can adjust the file size and the number of files to be retained in the log4j2.xml configuration file, as needed.

For more information about Log4j 2, see the Log4j 2 open-source project.

HTTP request logging

HTTP requests to the runtime engine and the administrative console are logged to the <date>.request.log file and <date>.request2.log, respectively, by the Pingfederate web container.

Like other PingFederate-generated log files, the HTTP request logs are written to the default PingFederate log directory. Properties controlling request logging are contained in the web-container configuration files:

- jetty-runtime.xml for the runtime engine (the <date>.request.log files)
- jetty-admin.xml for the administrative console (the <date>.request2.log files)

You can find these files in the <pf_install>/pingfederate/etc directory, and you can independently manage them on a per-server basis.

Administrator audit logging

PingFederate records actions performed by server administrators.

This information is recorded in the $<\!pf_install>$ /pingfederate/log/admin.log file. While the events themselves are not configurable, you can adjust Log4j 2 configuration settings to deliver the desired level of detail surrounding each event in the $<\!pf_install>$ /pingfederate/server/default/conf/log4j2.xml file.

Events logged by PingFederate include but are not limited to:

- Login attempt
- Explicit user logout (no time-outs)
- Account activation or deactivation
- Password change or reset
- Role change
- System settings management
- Certificate management
- OAuth settings management
- Metadata export
- XML file signatures applied
- Configuration archive export and import
- Identity provider (IdP)/service provider (SP) adapter, IdP token processor, or SP token generator created, modified, or deleted
- IdP/SP default URLs modified
- IdP/SP connection created, modified, or deleted
- Adapter-to-Adapter mapping or token exchange mapping created, modified, or deleted
- Authentication policy contract created, modified, or deleted
- IdP Discovery management
- SP Affiliation created, modified, or deleted
- PingOne for Enterprise account connected, modified, or disconnected

Each entry in the admin.log file is on a separate line and represents a single administrator action. The general format of each entry is the same, though specific events are recorded with information relevant to each type. Events are recorded when you click the corresponding **Save** button in the administrative console. Each log entry contains information relating to the event, including:

- The time the event occurred on the PingFederate server
- The username of the administrator performing the action

- The roles assigned to the administrator at the time the event occurred
- The type of event that occurred
- Basic information about the event

Each of these fields is separated by a vertical pipe (|) for easier parsing.

Detailed event logging

You can also configure PingFederate to log additional event information to a separate log file. When you enable detailed event logging, besides writing basic information to <pf install>/pingfederate/ log/adming.log, PingFederate logs detailed information about each event to admin-eventdetail.log in the same log directory. PingFederate links events between admin.log and adminevent-detail.log by a unique event ID. Each entry in the admin-event-detail.log file contains:

- The ID of the event
- The name of the file involved
- The type of event that occurred
- The line number where the change occurred
- The changes made



Not all events have detailed information. For example, login attempts are only logged to admin.log.

To enable detail event logging, set the pf.log.eventdetail property to true in the <pf install>/ pingfederate/bin/run.properties file.

API audit logging

PingFederate provides API endpoints and management services on the administrative port (9999) and the runtime port (9031) that are logged for auditing purposes.

Actions performed through these endpoints are logged for auditing purposes, as described in the following table.

API	Port	Log File
Administrative API	Administrative Port	admin-api.log
OAuth Client Management Service	Runtime Port	runtime-api.log
OAuth Access Grant Management Service	Runtime Port	runtime-api.log
Session Revocation API	Runtime Port	runtime-api.log

Administrative API audit log

PingFederate records actions performed through the administrative API in the <pf install>/ pingfederate/log/admin-api.log file.

While the events are not configurable, Log4j 2 configuration settings in the <pf install>/ pingfederate/server/default/conf/log4j2.xml file can be adjusted to deliver the desired level of detail surrounding each event.

Each log entry contains information relating to the event, including:

- Time the event occurred on the PingFederate server
- Administrator username performing the action
- Authentication method
- Client IP

- HTTP method
- REST endpoint
- HTTP status code

Each of these fields is separated by a vertical pipe (|) for ease of parsing.

Runtime APIs audit log

PingFederate records actions performed through the OAuth Client Management Service, the OAuth Access Grant Management Service, and the Session Revocation API in the cpf_install>/
pingfederate/log/runtime-api.log file.

While the events are not configurable, Log4j 2 configuration settings in the $<\!pf_install>/$ pingfederate/server/default/conf/log4j2.xml file can be adjusted to deliver the desired level of detail surrounding each event.

Each log entry contains information relating to the event, including:

- Time the event occurred on the PingFederate server
- Administrator username performing the action
- Authentication method
- Client IP
- HTTP method
- REST endpoint
- HTTP status code

Each of these fields is separated by a vertical pipe (|) for ease of parsing.

Runtime transaction logging

PingFederate provides for flexible, scalable logging of all federated-identity transactions, for both inbound and outbound messages.

About this task

Administrators can configure transaction logging to any of the four modes on a per-connection basis or override the logging mode for all service provider (SP) connections, identity provider (IdP) connections, or both for troubleshooting or as a one-step means of raising or lowering all connection logging modes to the same level. The log file is transaction.log, located in the $pf_install>>/pingfederate/log$ directory.

The following table describes the four transaction logging modes.

Mode	Description
None	No transaction logging.
Standard	(Default) Summary information for each transaction message, including:
	 Time stamp Hostname and port Log mode Connection ID SAML status code, for SAML responses only Context Message type SAML ID, for SAML messages only Endpoint, for outbound messages only Target URL, if single sign-on (SSO) transaction

Mode	Description
Enhanced	Includes everything logged at the Standard level including:
	 SAML_SUBJECT* Binding Relay state, if available Signature policy Signature status HTTP request parameters, for outbound messages only
	* Only when available in a SAML assertion, a single logout (SLO) request, an STS Request Security Token Response (RSTR), or an authentication request (AuthnRequest)
Full	Includes everything logged at the Enhanced level plus the complete XML message for every transaction.

Each field is separated by a vertical pipe (|) for parsing.

Steps

- To configure transaction logging mode on a per connection basis:
 - a. Select the applicable connection on the IdP Connections window (Authentication # Integration # IdP Connections) or the SP Connections window (Applications # Integration # SP Connections).
 - b. On the **General Info** tab, select one of the logging modes.
- To override transaction logging mode for all SP or IdP connections:
 - a. On the IdP Connections window or SP Connections window, click Show Advanced Fields.
 - b. On the **Logging Mode Override** setting, click **On**.
 - c. Select a logging mode for the IdP or SP connections.

Security audit logging

PingFederate records a subset of transaction log information with additional details at runtime, intended to facilitate security auditing and regulatory compliance.

The system records activities from single sign-on (SSO), single logout (SLO), OAuth, WS-Trust STS, and System for Cross-domain Identity Management (SCIM) inbound provisioning transactions in the security audit log, the audit.log file, located in the <pf install>/pingfederate/log directory. You can output security audit log information to different formats, including databases, CEF, and Splunk.



Note:

Outbound provisioning transactions are not included in the security audit log. Instead, they are recorded in the outbound provisioning audit log, the provisioner-audit.log file, located in the <pf install>/ pingfederate/log directory.

The following tables describe the default and available fields. PingFederate separates each field by a vertical pipe (|). As needed, fields are configurable by editing the <pf install>/pingfederate/ server/default/conf/log4j2.xml file.

Default fields, in the order that PingFederate records them in the security audit log

Field	Description
%d	The transaction time.

Field	Description
trackingid	The tracking ID values uniquely identify user sessions, useful for correlating log messages in the audit and server logs.
event	The type of transaction. For example, SSO, OAuth, AUTHN_ATTEMPT, AUTHN_REQUEST, AUTHN_SESSION_CREATED, AUTHN_SESSION_USED, AUTHN_SESSION_DELETED, and SRI_REVOKED.
	AUTHN_ATTEMPT and AUTHN_REQUEST indicate an authentication attempt against an identity provider (IdP) adapter instance and an authentication request sent to another IdP partner, through an IdP connection, respectively.
	AUTHN_SESSION_CREATED and AUTHN_SESSION_USED indicate the creation and employing of a PingFederate session, respectively.
	AUTHN_SESSION_DELETED indicates that a PingFederate session has been removed as a result of a front-channel browser-based logout request via the SAML 2.0 or WS-Federation protocol.
	SRI_REVOKED indicates that a PingFederate session has been added to the session revocation list.
subject	The subject of the transaction or authentication attempt.
ip	The incoming IP address.
арр	The target service provider (SP) application, the email verification endpoint, or the profile management page, when applicable and available.
connectionid	The partner identifier associated with the transaction. The OAuth client ID value for OAuth transactions. The ID of the authentication policy contract referenced by the local identity profile that has been invoked for the purpose of accessing the email verification endpoint or the profile management page.
protocol	The associated identity protocol; for example, SAML20 or OAuth20.
host	The host name or IP address of the PingFederate server.
role	The role PingFederate played for the transaction.
status	The status of the transactions.
adapterid	The ID of an adapter instance.
	Consider adding the authenticationsourceid and targetsessionid fields to record additional information about the request.
description	The description of an authentication failure, when such information is available from the authentication source, or an authorization failure from an erroneous OAuth authorization request.
responsetime	The time elapsed in milliseconds from when the system receives a final request for a transaction, to when the system writes the audit message. This value serves as an approximation of total transaction processing time and can be useful for monitoring trends.

Other available fields, in alphabetical order

Field	Description
accessgrantguid	The GUID of the OAuth access grant, for OAuth transactions.
assertionid	The unique ID for the SAML assertion.

Field	Description	
attrackingid	The tracking ID for OAuth access token. You can use this ID to analyze the flow of OAuth access tokens in the audit log and between PingFederate and PingAccess.	
attributes	The user attributes received (for an SP log), sent (for an IdP log), or provided by the user through the self-service registration or profile management page.	
authenticationsourceid	An array of one or more IdP adapters, one or more IdP connections, and identity profile, if any, invoked in an authentication or logout flow. For example, [adapter.HTMLFormSimplePCV, idpConnection.IdP, localIdentity.A8me9rySDn1aIM48]	
authnsessionexpiry	The expiry of an authentication session that has just been created or used.	
connectionname	The partner name associated with the transaction. The OAuth client name for OAuth transactions. The name of the authentication policy contract referenced by the local identity profile that has been invoked for the purpose of accessing the email verification endpoint or the profile management page.	
granttype	The OAuth grant type.	
header{ <i>anHttpRequestH</i>	eadles HTTP request header value identified by the header name. The header name is case-insensitive. For example, header {user-agent} and header {User-Agent} are equivalent.	
	To record multiple headers, repeat the header field, as illustrated in the following sample pattern.	
	<pre><pattern> %header{accept-language} %header{dnt} %n</pattern></pre> <pre>pattern></pre>	
	Given this partial sample, PingFederate includes both the accept-language and dnt HTTP request header values when recording entries in the audit log.	
	Note:	
	To record values from all HTTP request headers, look for the org.sourceid.servlet.filter.HttpRequestHeaderFilter Logger in the log4j2.xml file.	
	This capability is turned off by default and is likely suitable only for testing and troubleshooting purposes.	
httprequestid	The ID of the HTTP request. This can be used for correlation across external systems (like PingDirectory) and for debugging purposes in the server log. This field is optional.	
initiator	The federation role that initiated the SSO or SLO: SP or IDP.	
	Applicable only to SAML 2.0 transactions.	
inmessagetype	The incoming message type.	
	Possible values are Request or Response.	
inresponseto	The value of the InResponseTo attribute of an SSO or SLO response.	
inxmlmsg	The incoming message. For example, a SAML AuthnRequest or the information pertaining to an OAuth request.	

Field	Description	
localuserid	The local ID used for the transaction, when account linking is enabled at the SP.	
outurl	The URL where the protocol response was sent. For security reason, parameters and fragments are excluded.	
outxmlmsg	The outgoing message. For example, a SAML Response or the information pertaining to a response for an OAuth request.	
parameter{ <i>anHttpReque</i>	st Panarretus)of the HTTP request parameter identified by the parameter name. The parameter name is case-sensitive.	
	To record multiple parameters, repeat the parameter field, as illustrated in the following sample pattern.	
	<pre><pattern> %parameter{foo1} %parameter{Foo3} %n</pattern></pre>	
	Given this partial sample, PingFederate includes both the foo1 and Foo3 HTTP request parameter values when recording entries in the audit log.	
	Note:	
	To record values from all HTTP request parameter, look for the org.sourceid.servlet.filter.HttpRequestParameterFilter Logger in the log4j2.xml file.	
	This capability is turned off by default and is likely suitable only for testing and troubleshooting purposes.	
pfversion	The PingFederate version.	
requestid	The ID of a SAML request.	
requeststarttime	The start time of the request in milliseconds since midnight, January 1, 1970 UTC.	
responseid	The ID of a SAML response.	
sessiongroupid	The internal ID for a group of persistent authentication sessions associated with a single browser instance through the PF.PERSISTENT cookie. It is only set if the request has triggered a session lookup.	
sri	The session reference identifier (SRI) for the user, which can be passed to the session revocation API to revoke the user's sessions. It is only set if the request has triggered a session lookup.	
stspluginid	The ID for the token processor or token generator instance.	
	Applicable only to WS-Trust STS transactions.	
targetsessionid	An array of one or more SP adapters or SP connections invoked in an authentication or logout flow.	

Field	Description	
trackedparameter{anHttp	preservation and is case-sensitive.	
	Tip:	
	The PingFederate policy engine is capable of tracking HTTP request parameters that it receives from the initial request and making them available to authentication sources, selector instances, and contract mappings throughout the policy. As needed, parameters can be configured as such on the Tracked HTTP Parameters window. For more information about tracked parameters, see <i>Policies</i> on page 239.	
	To record multiple parameters, repeat the trackedparameter field, as illustrated in the following sample pattern.	
	<pre><pattern> %trackedparameter{foo2} %trackedparameter{Foo4} %n</pattern></pre>	
	Given this partial sample, PingFederate includes both the foo2 and Foo4 HTTP request parameter values when recording entries in the audit log.	
	If the parameter, as indicated by <anhttprequestparameter>, has not been configured as a parameter to be tracked by the policy engine, PingFederate does not record the parameter value in the audit log.</anhttprequestparameter>	
validatorid	The ID of the Password Credential Validator (PCV) instance, for the successful attempts.	
virtualserverid	The virtual server ID of a request, if applicable.	

Outbound provisioning audit logging

The PingFederate provisioner-audit.log file records outbound provisioning transactions, intended to facilitate security auditing.

The provisioner-audit.log log file is located in the $< pf_install > / pingfederate/log$ directory. Outbound provisioning audit log information can be output to different formats, including database and Splunk.

The following table describes all recorded elements. Optionally, you can configure elements by editing the $< pf_install > /pingfederate/server/default/conf/log4j2.xml file.$

Item	Description	
%d	Transaction time.	
cycle_id	The unique ID for each provisioning cycle.	
channel_id	The unique ID of the provisioning channel between source and target.	
event_type	The type of provisioning events, such as CREATE and UPDATE.	
source_id	The provisioning Source ID.	
target_id	The provisioning Target ID.	
is_success	A flag to show whether the event was successful or not. If the attempt succeeded, the value is true; otherwise, the value is false.	
non_success_ cause	Description of failure cause.	

When PingFederate is configured to log DEBUG messages for troubleshooting purposes, it records all runtime and administrative events that can be used for troubleshooting in the <pf install>/ pingfederate/log/server.log file, including status and error messages.

Server log information can be output to a database server.



DEBUG messages are turned off by default. For troubleshooting purpose, you can re-enable it by editing the <pf install>/pingfederate/server/default/conf/log4j2.xml file.

The following table describes the recorded elements. Optionally, you can configure elements by editing the log4j2.xml file as well.

Item	Description
%d	Event date and time.
%X{trackingid}	The tracking ID values uniquely identify user sessions, useful for correlating log messages in the audit and server logs.
%p	Logging level.
%с	The Java class issuing the status or error message, when applicable.
%m	Status or error message.

To facilitate troubleshooting, administrators can use a filter utility to aggregate related events using the log filter tool.

Server log filter

PingFederate provides a utility, logfilter, that administrators can use to filter server logs.

The logfilter utility is located in the <pf install>/pingfederate/bin directory, logfilter.bat for Windows, and logfilter.sh for Linux.

The utility sorts through all the server logs in the log directory. Administrators can move or copy one or more server log files to a different directory that can be specified as an input parameter.

The log filter returns lists of log entries based on either:

- Entity ID and subject
- Tracking ID
- Session cross-reference ID

The following table describes the utility's command options. The table afterward describes optional parameters available for all of the commands.

Server log filter command parameters

Command parameter	Description
-entityid entity ID	These two commands must be used together and return a list of transactions
-subject subject	for the specified federation partner's entity ID and transaction subject.
-trackingid tracking	This command returns a list of transactions with the same tracking ID.

Command parameter	Description	
-sessionxrefid session cross-reference ID	This command returns a list of transactions for an ID assigned by PingFederate to associate different transactions according to the user session under which they occurred. The value of session cross-reference ID can be the value of any of the following transaction tags in the target server logs:	
	ArtifactSession IndexAssertion ID	

Server log filter parameters (optional)

Parameter	Description
-logsdir log files directory	Full or relative path to source directory for the logs.
	Default: all server.log files are written to the <pf_install>/ pingfederate/log directory, a setting that can be adjusted by the pf.log.dir property in the <pf_install>/pingfederate/bin/ run.properties file.</pf_install></pf_install>
-outputfile output file	Output path and file for the returned list.
	Default: \$ <pf.log.dir>/logfilter_output.log.</pf.log.dir>
-outputtoconsole	Returns list to the command console rather than to a file.

The log filter creates its own log file, logfilter.log, located in the log directory. Optionally, administrators can control settings for this log in the <pf install>/pingfederate/bin/ logfilter.log4j2.xml file.

Logging in other formats

PingFederate provides the option of writing the audit log, the provisioner audit log, the provisioner log, and the server log to commonly used databases with failover to file logging.

For the audit log and the provisioner audit log, administrators can choose instead to write the information to the Common Event Format (CEF), a differently formatted log file that can be used by Splunk, or both.

Writing logs to databases

Database logging replaces file logging. For each qualified database server, PingFederate provides scripts to create database tables for the audit log, the provisioner audit log, the provisioner log, and the server log.

About this task

You can find these scripts in the <pf install>/pingfederate/server/default/conf/log4j/ sql-scripts directory.



Note:

PingFederate was tested with vendor-specific JDBC drivers. For more information, see *Database driver* information on page 115. To obtain the database driver JAR file, contact your database vendor. Install the database driver file to the <pf install>/pingfederate/server/default/lib directory then restart the server.

Failover file logging is provided in the event that database logging fails for any reasons. By default, PingFederate retries database logging every minute. Messages written to log files during failover periods are not copied over to the database server.

You enable database logging for the audit log, the provisioner audit log, the provisioner log, and the server log in the log4j2.xml file.

Steps

- 1. Edit <pf install>/pingfederate/server/default/conf/log4j2.xml.
- 2. After the Preserve messages in a local file section, for each log that you want to enable database logging, uncomment the preset Java Database Connectivity (JDBC) appender configuration based on the choice of your database server.

Audit log

- Oracle MySQL SecurityAuditToMySQLDB
- Oracle Database SecurityAuditToOracleDB
- PostgreSQL SecurityAuditToPostgreSQLDB
- Microsoft SQL Server SecurityAuditToSQLServerDB

Provisioner audit log

- Oracle MySQL OutboundProvisionerEventToMySQLDB
- Oracle Database OutboundProvisionerEventToOracleDB
- PostgreSQL OutboundProvisionerEventToPostgreSQLDB
- Microsoft SQL Server OutboundProvisionerEventToSOLServerDB

Provisioner log

- Oracle MySQL ProvisionerLogToMySQLDB
- Oracle Database ProvisionerLogToOracleDB
- PostgreSQL ProvisionerLogToPostgreSQLDB
- Microsoft SQL Server ProvisionerLogToSOLServerDB

Server log

- Oracle MySQL ServerLogToMySQLDB
- Oracle Database ServerLogToOracleDB
- PostgreSQL ServerLogToPostgreSQLDB
- Microsoft SQL Server ServerLogToSQLServerDB



Note:

Each JDBC appender is followed by two related appenders, PingFailover and RollingFile. Together, they create a running *-failover.log file in the log directory in the event that database logging fails for any reason. Both appenders must also be enabled (uncommented).



For more information about each appender, review inline comments and notes in the log4j2.xml file.

3. Replace placeholder parameter values in log4j2.db.properties in the same conf directory for the applicable Java Database Connectivity (JDBC) servers.

The parameter values provide access to the database. Test and validate access prior to production deployment. Like log4j2.xml, log4j2.db.properties is also individually managed per

PingFederate server. This flexibility allows multiple PingFederate nodes in a clustered environment to write messages to different destinations, as needed.



Tip:

You can obfuscate the password used to access the database by running the obfuscate utility, located in the <pf install>/pingfederate/bin directory: obfuscate.bat for Windows or obfuscate.sh for Linux. Use the actual password as an argument and copy the entire result into the value for the password parameter in log4j2.db.properties.

4. Uncomment the appender reference, <appenderRef/>, in the associated logger elements, as described inline in the log4j2.xml file.

Audit log

Uncomment the corresponding PingFailover appender references from the following Logger elements located under the Loggers section:

- Browser SSO SP and adapter-to-adapter org.sourceid.websso.profiles.sp.SpAuditLogger
- Browser SSO IdP and adapter-to-adapter org.sourceid.websso.profiles.idp.IdpAuditLogger
- OAuth authorization server org.sourceid.websso.profiles.idp.AsAuditLogger
- Dynamic Client Registration org.sourceid.websso.profiles.idp.ClientRegistrationAuditLogger
- WS-Trust STS, IdP, and SP org.sourceid.wstrust.log.STSAuditLogger

Provisioner audit log

Uncomment the corresponding PingFailover appender reference from the ProvisionerAuditLogger Logger element located under the Set up the Outbound provisioner audit logger section.

Provisioner loa

Uncomment the corresponding PingFailover appender reference from the com.pingidentity.provisioner AsyncLogger element located under the Loggers section.

Server log

Uncomment the corresponding PingFailover appender reference from the root element located under the Set up the Root Logger section, near the end of the file.



Important:

As indicated in the IMPORTANT comments for the loggers, you must also remove some of the existing appender references.

5. Optional: For the audit log and the provisioner audit log, you can configure elements for database logging in the ConversionPattern appender parameter, as needed.

Logging in Common Event Format

You can use PingFederate to write from logs using the Common Event Format (CEF) open standard.

CEF is an open logging standard. PingFederate provides an option of writing elements from the audit log, the provisioner audit log, or both at runtime to a syslog receiver for parsing and analysis using ArcSight from Micro Focus. Alternatively, administrators can write the information to a flat file in CEF. However, you should use syslog when available.



PingFederate is tested with ArcSight for interoperability using the default elements defined in log4j2.xml. Any additions to these elements might render your CEF logging incompatible with ArcSight.

Writing audit log in CEF

You can write the audit log in Common Event Format (CEF) in PingFederate.

Steps

- 1. Edit cpf install>/pingfederate/server/default/conf/log4j2.xml.
- 2. Under the Security Audit log: CEF Formatted syslog appender section, uncomment one of the preset appender configurations:
 - SecurityAuditToCEFSyslog a Socket appender
 - SecurityAuditToCEFFile a RollingFile appender



The SecurityAuditToCEFSyslog Socket appender is followed by two related appenders, PingFailover and RollingFile. Together, they create a running audit-cef-syslogfailover.log file in the log directory in the event that CEF logging fails for any reason. Both appenders must also be enabled and uncommented.



Review inline comments and notes in the log4j2.xml file for more information about each appender.

- 3. If you are configuring the SecurityAuditToCEFSyslog Socket appender, replace the placeholder parameter values for the syslog host.
- 4. If you are configuring the SecurityAuditToCEFSyslog Socket appender. uncomment the PingFailover appender reference (<appender-ref ref="SecurityAuditToCEFSyslog-FAILOVER"/>) from the following Logger elements located under the Loggers section:
 - Browser SSO SP and adapter-to-adapter org.sourceid.websso.profiles.sp.SpAuditLogger
 - Browser SSO IdP and adapter-to-adapter org.sourceid.websso.profiles.idp.IdpAuditLogger
 - OAuth authorization server org.sourceid.websso.profiles.idp.AsAuditLogger
 - Dynamic Client Registration org.sourceid.websso.profiles.idp.ClientRegistrationAuditLogger
 - WS-Trust STS, identity provider (IdP), and service provider (SP) org.sourceid.wstrust.log.STSAuditLogger



Important:

As indicated in the IMPORTANT comments for the loggers, you must also remove some of the existing appender references.

Writing provisioner audit log in CEF

You can write provisioner audit logs in Common Event Format (CEF) for PingFederate. PingFederate provides an option of writing elements from the audit log and the provisioner audit log at runtime to a syslog receiver for parsing and analysis using ArcSight from Micro Focus.

Steps

- 1. Edit <pf install>/pingfederate/server/default/conf/log4j2.xml.
- 2. Uncomment one of the preset appender configurations:
 - OutboundProvisionerEventToCEFSyslog (a Socket appender under the Outbound provisioner audit log: CEF Formatted syslog appender section)



Note:

This Socket appender is followed by two related appenders, PingFailover and RollingFile. Together, they create a running provisioner-audit-cef-syslogfailover.log file in the log directory in the event that CEF logging fails for any reason. Both appenders must also be enabled (uncommented).

• OutboundProvisionerEventToCEFFile (a RollingFile appender under the Outbound provisioner audit log for CEFFile section)



Z Tip:

Review inline comments and notes in the log4j2.xml file for more information about each appender.

- 3. If you are configuring the OutboundProvisionerEventToCEFSyslog Socket appender, replace the placeholder parameter values for the syslog host.
- 4. If you are configuring the OutboundProvisionerEventToCEFSyslog Socket appender, uncomment the PingFailover appender reference (<appender-ref ref="OutboundProvisionerEventToCEFSyslog-FAILOVER"/>) from the ProvisionerAuditLogger Logger elements located under the Set up the Outbound provisioner audit logger section.



Important:

As indicated in the IMPORTANT comments for the loggers, you must also remove some of the existing appender references.

Writing audit log for Splunk

Ping Identity provides a custom Splunk App for PingFederate to process audit logs generated by a PingFederate deployment. Splunk is an enterprise software that allows for monitoring, reporting, and analysis of consolidated log files.

Before you begin

Download and install Splunk

About this task

Splunk captures and indexes real-time data into a single searchable repository where reports, graphs, and other data visualization can be generated.

The PingFederate Splunk App provides rich system monitoring and reporting, including:

Current transaction and system reports

- Service reports, such as a daily usage report, and identity provider (IdP) and service provider (SP) reports per connection
- Trend reports, such as weekly and monthly usage reports, and trend analysis

Splunk uses a specially formatted version of the audit log splunk-audit.log, which you can write to the PingFederate log directory when you complete the setup steps.



Note:

The Splunk App for PingFederate is available separately. It requires enterprise-licensed, or trial installation of the Splunk software and the Splunk Universal Forwarder, which is needed to collect data from the PingFederate audit log for Splunk. The application includes additional documentation on installation and available features. To download the free application, go to splunkbase.splunk.com and search for PingFederate.

Steps

- 1. Set up your Splunk server.
 - a. Enable a receiver to listen for data from the PingFedrate server.

For more information, see Splunk documentation.

- b. Install Splunk App for PingFederate.
- 2. Configure PingFederate to write audit log messages to the <pf install>/pingfederate/log/ splunk-audit.log file.
 - a. Edit <pf install>/pingfederate/server/default/conf/log4j2.xml.
 - b. Locate the following Logger elements located under the Loggers section:
 - Browser single sign-on (SSO) SP and adapter-to-adapter org.sourceid.websso.profiles.sp.SpAuditLogger
 - Browser SSO IdP and adapter-to-adapter org.sourceid.websso.profiles.idp.IdpAuditLogger
 - OAuth authorization server org.sourceid.websso.profiles.idp.AsAuditLogger
 - Dynamic Client Registration org.sourceid.websso.profiles.idp.ClientRegistrationAuditLogger
 - WS-Trust STS, IdP, and SP org.sourceid.wstrust.log.STSAuditLogger
 - c. Uncomment the SecurityAudit2Splunk RollingFile appender reference, <appenderref ref="SecurityAudit2Splunk"/>, from one or more of the Logger elements.

Example:

For example, the default logger for an IdP audit log reads as follows.

```
<Logger name="org.sourceid.websso.profiles.idp.IdpAuditLogger"</pre>
        level="INFO" additivity="false" includeLocation="false">
    <appender-ref ref="SecurityAudit2File" />
        <appender-ref ref="SecurityAuditToCEFSyslog-FAILOVER"/>
        <appender-ref ref="SecurityAuditToCEFFile"/>
        <appender-ref ref="SecurityAuditToMySQLDB-FAILOVER"/>
        <appender-ref ref="SecurityAuditToPostgreSQLDB-FAILOVER" />
        <appender-ref ref="SecurityAuditToSQLServerDB-FAILOVER"/>
        <appender-ref ref="SecurityAuditToOracleDB-FAILOVER"/>
        <appender-ref ref="SecurityAudit2Splunk"/>
```

```
</Logger>
```

To log Browser SSO IdP audit log messages to splunk-audit.log, update the Logger element as follows.

```
<Logger name="org.sourceid.websso.profiles.idp.IdpAuditLogger"</pre>
        level="INFO" additivity="false" includeLocation="false">
    <appender-ref ref="SecurityAudit2Splunk"/>
    <!--
        <appender-ref ref="SecurityAuditToCEFSyslog-FAILOVER"/>
        <appender-ref ref="SecurityAuditToCEFFile"/>
        <appender-ref ref="SecurityAuditToMySQLDB-FAILOVER"/>
        <appender-ref ref="SecurityAuditToPostgreSQLDB-FAILOVER" />
        <appender-ref ref="SecurityAuditToSQLServerDB-FAILOVER"/>
        <appender-ref ref="SecurityAuditToOracleDB-FAILOVER"/>
        <appender-ref ref="SecurityAudit2Splunk"/>
        <appender-ref ref="SecurityAudit2File" />
</Logger>
```

Note:

For auditing of adapter-to-adapter events, you must enable both the IdP and SP loggers.

d. Uncomment the following section:

```
<RollingFile name="SecurityAudit2Splunk" fileName="${sys:pf.log.dir}/</pre>
splunk-audit.log"
filePattern="${sys:pf.log.dir}/splunk-audit.%d
{yyyy-MM-dd}
.log"
ignoreExceptions="false">
<PatternLayout>
<pattern>%d trackingid="%X
{trackingid}
" event=%X
{event}
subject="%X
{subject}
" ip=%X
{ip}
app=%X
connectionid=%X
{connectionid}
protocol="%X
{protocol}
" pfhost=%X
{host}
role=%X
```

```
{role}
status=%X
{status}
adapterid=%X
{adapterid}
description="%X
{description}
" responsetime=%X
{responsetime}
inmessagetype="%X
{inmessagetype}
" %n</pattern>
</PatternLayout>
<Policies>
<TimeBasedTriggeringPolicy />
</Policies>
</RollingFile>
```

- 3. Set up Splunk Universal Forwarder.
 - a. Download the Splunk Universal Forwarder from Splunk and install it on the PingFederate server.
 - b. Configure the Splunk Universal Forwarder to monitor the splunk-audit.log file and forward the data to the receiver configured in *step 1b*.

For detailed installation and configuration instructions, see Splunk documentation.

Alternative console authentication

As an alternative to using PingFederate's own internal datastore for authentication to the administrative console, you can configure PingFederate to use your network's LDAP user-datastore, the RADIUS protocol, client certificates, or OIDC-based authentication.

You can configure any of these alternative console authentication methods at any time. Most user-management functions are handled outside the scope of the PingFederate administrative console when alternative authentication is enabled.

Unlike native authentication, for which you configure local accounts and their privileges in **System** # **Server** # **Administrative Accounts**, you must define roles in configuration files when using an alternative authentication scheme. Similar to native authentication, PingFederate provides two account types and three administrative roles for role-based access control, as shown in the following table.

PingFederate User Access Control

Account type	Administrative role	Access privileges
Admin	User Admin	Create users, deactivate users, change or reset passwords, and install replacement license keys.
Admin	Admin	Configure partner connections and most system settings, except the management of local accounts and the handling of local keys and certificates.

Account type	Administrative role	Access privileges
	Expression Admin	Map user attributes by using the expression language, Object-Graph Navigation Language (OGNL).
		Important:
		Only Administrative users who have both the Admin role and the Expression Admin role:
		 Can be granted the User Admin role. This restriction prevents non-Expression Admin users from granting themselves the Expression Admin Role.
		 Can be granted write access to the file system or directory where PingFederate is installed. This restriction prevents a non- Expression Admin user from placing a data.zip file containing expressions into the <pf_install>/pingfederate/server/ default/deploy directory, which would introduce expressions into PingFederate.</pf_install>
Admin	Crypto Admin	Manage local keys and certificates.
Auditor	Not applicable	View-only permissions for all administrative functions. When the Auditor role is assigned, no other administrative roles can be set.



All four administrative roles are required to access and make changes through the following services:

- The /bulk, /configArchive, and /configStore administrative API endpoints
- The Configuration Archive window, accessed from System # Server, in the administrative console
- The Connection Management configuration item on the Service Authentication window, accessed from Security # System Integration

Enabling LDAP authentication

You can enable LDAP authentication by using the configuration files located in the <pf install>/ pingfederate/bin directory.

About this task

When LDAP authentication is configured, PingFederate does not lock out administrative users based upon the number of failed sign-on attempts. Instead, responsibility for preventing access is delegated to the LDAP server and enforced according to its password lockout settings.

Steps

1. In the <pf install>/pingfederate/bin/run.properties file, change the value of the pf.console.authentication property as shown.

pf.console.authentication=LDAP

2. In the <pf install>/pingfederate/bin/ldap.properties file, change property values as needed for your network configuration.

For more information, see the comments in the file.

The roles configured in the properties file apply to both the administrative console and the administrative API.



Important:

Remember to assign LDAP users or designated LDAP groups to at least one of the PingFederate administrative roles as indicated in the properties file.



You can also use this configuration file in conjunction with RADIUS authentication to determine permissions dynamically through an LDAP connection.

3. Start or restart PingFederate.

Enabling RADIUS authentication

You can enable RADIUS authentication using the configuration files located in the <pf install>/ pingfederate/bin directory. The RADIUS protocol provides a common approach for implementing strong authentication in a client-server configuration.

About this task

PingFederate supports the protocol scenarios for one-step authentication, such as appending a one-time passcode obtained from an authenticator to the password, and two-step authentication, such as through a challenge-response process.



Note:

When RADIUS authentication is configured, PingFederate does not lock out administrative users based on the number of failed sign-on attempts. Instead, responsibility for preventing access is delegated to the RADIUS server and enforced according to its password lockout settings.



The NAS-IP-Address attribute is added to all Access-Request packets sent to the RADIUS server. The value is copied from the pf.engine.bind.address property in run.properties. Only IPv4 addresses are supported.

Steps

1. In the <pf install>/pingfederate/bin/run.properties file, change the value of the pf.console.authentication property as shown.

pf.console.authentication=RADIUS

2. In the <pf install>/pingfederate/bin/radius.properties file, change property values as needed for your network configuration.

For more information, see the comments in the file.

The roles configured in the properties file apply to both the administrative console and the administrative API.



Important:

Be sure to assign RADIUS users or designated RADIUS groups to at least one of the PingFederate administrative roles as indicated in the properties file. Alternatively, you can set the use.ldap.roles property to true and use the LDAP properties file, also in the bin directory, to map LDAP groupbased permissions to PingFederate roles.

3. Start or restart PingFederate.

Multi-factor console authentication using PingID

PingID is a cloud service that enables multi-factor authentication (MFA) using a mobile application.

The PingFederate administrative console supports authentication through the RADIUS protocol, which provides a common approach for implementing strong authentication in a client-server configuration.

By combining these two capabilities, you can configure PingID to provide MFA to protect access to the PingFederate administrative console, which meets the requirement of stronger authentication for administrators accessing security-related software products.



PingID requires a separate license. Please contact sales@pingidentity.com or request a trial license at pingidentity.com.

Requirements

The following components are required to enable MFA to the administrative console using PingID:

- PingFederate with external access to the PingID cloud service
- A PingID license
- A directory server where the administrative credentials and group membership are stored

Solution overview

You can enable multi-factor console authentication using PingID.

When PingID is the second authentication factor for the PingFederate administrative console, the administrators must authenticate successfully against the first factor, such as a directory server, and subsequently respond to the request for authentication from the PingID app on their mobile devices.

Multi-factor Console Authentication using PingID

Processing steps

- 1. An administrator opens a browser and accesses the PingFederate administration console.
 - a. The administrative console displays the Sign On page.
 - b. The administrator enters the correct username and password.
- 2. PingFederate invokes the PingID Password Credential Validator (PCV) to validate the username and password against your directory server.
 - The PingID PCV comes with a built-in RADIUS server, which can be used as the point of authentication for the PingFederate administration console using RADIUS authentication.
- 3. Upon successful validation of the user credentials, the PingID PCV invokes the PingID service with the username.

The PingID service looks for the username in its datastore.

- If the administrator has not registered a device for use with PingID, the PingID service returns a "username unknown" message. The administrative console displays a device registration window. The administrator must register the mobile device.
- If the administrator has a registered device, the PingID service notifies the PingID app on the device or sends a text message (SMS) or voice callback message, depending on the configuration for that user account.
 - a. The administrator responds to the request for authentication from PingID.
 - b. If the administrator has successfully authenticated to the PingID notification, the PingID service returns a "success" message to the PingID PCV.
- 5. The administrative console menu opens.

Configuring your PingID account

To use PingID as the second factor to authenticate into the PingFederate administrative console, you need a PingID or PingOne account. You must also configure the account to allow the PingID Password Credential Validator (PCV) to use the PingID service.

Steps

- 1. Register a PingID account.
 - a. Contact sales@pingidentity.com for a registration key.
 - b. Register at *PingOne admin portal*.
 - c. Select PingOne for Enterprise and then click Next.
 - Let the PingOne admin portal guide you through the registration process. At the end, a confirmation email is sent to your email address.
 - d. Open the confirmation email and follow the instructions to activate your PingOne for Enterprise account.
- 2. Enable PingID client integration.
 - a. Sign on to the PingOne admin portal.
 - b. Go to Setup # PingID # Client Integration.
 - c. Under Integrate with PingFederate and other Clients, click Download and save the pingid.properties file for a subsequent task.
 - d. Sign off of the PingOne admin portal.

Creating an LDAP Username Password Credential Validator instance

You can create an LDAP username password credential validator (PCV) in the PingFederate administrative console to create a second factor for multi-factor authentication (MFA).

About this task

Administrators must authenticate successfully against the first factor, such as a directory server where the administrator accounts, credentials and group memberships are stored. To fulfill this requirement, you need an LDAP connection from PingFederate to your directory server, and an instance of the LDAP Username Password Credential Validator.

Steps

- 1. Go to System # Data & Credential Stores # Password Credential Validators. On the Password Credential Validators window, click Create New Instance.
- 2. On the **Type** tab, from the **Type** list, select the **LDAP Username Password Credential Validator** and complete the **Instance Name** and **Instance ID** fields.
- 3. On the **Instance Configuration** tab, from the **LDAP datastore** list, select the datastore and complete the **Search Base** and **Search Filter** fields.

For more information about each field, see the following table.

Field	Description
LDAP Datastore	The LDAP datastore configured in PingFederate.
(Required)	If you have not configured the server to communicate with the LDAP directory server you need, click Manage Data Stores .
	There is no default selection.
Search Base	The location in the directory server where the search begins.
(Required)	This field has no default value.

Field	Description
Search Filter	The LDAP query to locate a user record.
(Required)	If your use case requires the flexibility of allowing users to identify themselves using different attributes, you can include these attributes in your query. For instance, the following search filter allows users to sign on using either the samaccountName or employeeNumber attribute value through the HTML Form Adapter:
	<pre>((sAMAccountName=\${username}) (employeeNumber= \${username}))</pre>
	This field has no default value.
Scope of Search	The level of search to be performed in the search base.
	One Level indicates a search of objects immediately subordinate to the base object, not including the base object itself. Subtree indicates a search of the base object and the entire subtree within the base object distinguished name.
	The default selection is Subtree .
Case-Sensitive Matching	The option to enable case-sensitive matching between the LDAP error messages returned from the directory server and the Match Expression values specified on this window.
	This check box is selected by default.

- 4. On the Extended Contract tab, click Next to skip to the Summary tab.
- 5. On the **Summary** tab, review the configuration, modify as needed, and then save the configuration.

Configuring a PingID Password Credential Validator instance

You can create and configure an instance of the PingID Password Credential Validator in the PingFederate administrative console.

Before you begin

Open the previously downloaded pingid.properties file in a text editor, copy its content, and then close the file.

Steps

- 1. Go to System # Data & Credential Stores # Password Credential Validators, and then click Create New Instance.
- 2. On the **Type** tab, from the **Type** list, select **PingID PCV** (with integrated **RADIUS server**) and complete the **Instance Name** and **Instance ID** fields.

- 3. On the **Instance Configuration** tab, configure the required fields as follows.
 - a. Click Add a new row to 'RADIUS Clients', enter 127.0.0.1 as the RADIUS IP address and a value in the Client Shared Secret field. Click Update.



Tip:

127.0.1 represents the local RADIUS client, the PingFederate administrative console, which calls the RADIUS server bundled in the PingID PCV for authentication.

The **Client Shared Secret** value is required for the next task.

- b. Click Add a new row to 'Delegate PCV's', select the previously created LDAP Username Credential Validator instance, and then click Update.
- c. Paste the content from the pingid.properties file into the PingID Properties File field.
- d. Review the rest of the default settings. Modify as needed to meet your requirements, and click Next.
- 4. On the **Extended Contract** tab, click **Next** to skip to the **Summary** tab.
- 5. On the **Summary** tab, review, modify if needed, and save the configuration.

Configuring PingFederate to use RADIUS authentication

You can enable RADIUS authentication in the PingFederate administrative console.

About this task

In this multi-factor console authentication use case, the PingFederate administrative console is a RADIUS client that calls the local RADIUS server bundled in the PingID Password Credential Validator (PCV) for the second factor authentication.



Note:

For a clustered PingFederate environment, perform these steps on the console node.

Steps

1. Open the <pf install>/pingfederate/bin/run.properties file in a text editor and set the pf.console.authentication property to RADIUS.

pf.console.authentication=RADIUS

2. Obfuscate the **Client Shared Secret** value using a PingFederate command-line tool.

Example:

Windows: <pf install>\pingfederate\bin\obfuscate.bat clientSharedSecret

Linux: <pf install>/pingfederate/bin/obfuscate.sh clientSharedSecret

Result:

The output should be a long line of text.

- 3. Copy the output for the next step.
- 4. Open the <pf install>/pingfederate/bin/radius.properties file in a text editor and modify as follows.

host=localhost shared.secret=obfuscatedClientSharedSecret



The timeout value is the number of milliseconds to wait for the second authentication factor to complete before timing out the login attempt. In this use case, ten seconds, or 10000 ms, should be sufficient for PingID.

In addition, assign one or more RADIUS users or designated RADIUS groups to at least one of the PingFederate administrative roles as indicated in the radius.properties file. Alternatively, you can set the use.ldap.roles property to true and use the LDAP properties file, ldap.properties in the same bin directory, to map LDAP group-based permissions to PingFederate roles.

5. Save your changes, and restart PingFederate.

Verifying your setup

Verify your setup after you have completed the required configuration.

Steps

- Start PingFederate.
 - In a clustered PingFederate environment, start PingFederate on the console node.
- 2. Start a web browser.
- 3. Browse to the URL.
 - https://*<pf_host>*:9999/pingfederate/app where *<pf_host>* is the network address of your PingFederate server. It can be an IP address, a host name, or a fully qualified domain name. It must be accessible from your computer.
- 4. Authenticate using your directory user credentials.
 - Upon successful validation, the PingFederate administrative console prompts you to authenticate using PingID if you have a registered device. If you do not have a registered device, follow the steps shown on-screen to register a device.
- 5. Respond to the authentication request from PingID, and the PingFederate administrative console menu will display.

Result



Important:

Access to the PingFederate administrative console can now only be performed by using directory user credentials, followed by a PingID authentication.

Enabling certificate-based authentication

You can enable certificate-based authentication in the PingFederate administrative console.

Before you begin

- Have a PingFederate username and password.
- Import the necessary client key and certificate into the web browser you use to access PingFederate.

About this task

To enable client-certificate authentication, PingFederate administrative users must import an X.509 key and a suitable certificate for user authentication into their web browsers. In addition, the corresponding root certificate authority (CA) certificates must be contained in the Java runtime or the PingFederate

trusted store. Other setup steps, including designating user permissions, must be completed by using configuration files located in the <pf install>/pingfederate/bin directory.

The roles configured in the properties file apply to both the administrative console and the administrative API.

Steps

- 1. Sign on to the PingFederate console as a user with permissions that include the Crypto Admin role.
- 2. Ensure the client-certificate's root CA and any intermediate CA certificates are contained in the trusted store, either for the Java runtime or PingFederate.



Note:

You can import a certificate to PingFederate in Security # Certificate & Key Management # Trusted CAs.



You might want to click the Serial Number and copy the Issuer distinguished name (DN) to use in later steps.

3. In the <pf install>/pingfederate/bin/run.properties file, change the value of the pf.console.authentication property as shown.

pf.console.authentication=cert

4. In the <pf install>/pingfederate/bin/cert auth.properties file, enter the Issuer DN for the client certificate as a value for the property rootca.issuer.x, where x is a sequential number starting at 1.



Note:

If you copied the Issuer DN after step 2, paste this value. For more information, see the comments in the file.

The roles configured in the properties file apply to both the administrative console and the administrative API.

- 5. Repeat the previous step for any additional CAs as needed.
- 6. Enter the certificate user's Subject DN for the applicable PingFederate permission roles, as described in the properties file.



Important:

The configuration values are case-sensitive.

7. Repeat the previous step for all users as needed.



Note:

Other settings in the properties file are used to display the user's ID (Subject DN) in abbreviated form in the administrative console.

8. Start or restart PingFederate.

The intent of automatic multi-connection error checking is to verify that all configured connections have not been adversely affected by subsequent changes in supporting components.

About this task

Automatic multi-connection error checking occurs when you access certain supporting components, such as the SP Adapters and IdP Adapters windows, the Token Processors window, the Token Generators window, the Password Credential Validators window, and the Identity Store Provisioners window.

As the number of connections and supporting components increases, so does the validation time. If you experience noticeable delays in accessing adapters, token translators, password credential validators, or identity store provisioner, you can turn off automatic connection validation.



Note:

When automatic connection validation is turned off, error checking is deferred until you access the connection lists on the administrative console. You can make configuration changes without being prompted to fix all dependency errors immediately, but remember that the configuration changes can potentially cause service disruption.

For example, if you remove a configured source-attribute in an attribute fulfillment configuration in a connection, users will be unable to complete single sign-on (SSO) requests until you reconfigure such connections.

When you access the Connections window, if the administrative console detects one or more dependency error conditions, it displays a visual cue to indicate the errors. To resolve each error, select the applicable connection and follow the on-screen instructions to modify the configuration.

This setting does not affect the validation of a connection being configured or modified. Also, individual connections are always validated automatically when you access them on the Connections window, regardless of the configuration of this setting.

To manage this setting for SP, go to Applications # Integration # SP Connections on the SP Connections window.

To manage this setting for IdP, go to Authentication # Integration # IdP Connections on the IdP Connections window.

Steps

 To turn off automatic multi-connection error checking, click Show Advanced Fields and select the Disable Automatic Connection Validation check box.

This check box is not selected by default.

Result:

After you select or clear the check box, the state of this setting is reflected on both the SP Connections window and the IdP Connections window.

Automating configuration migration

PingFederate provides a configuration-migration tool for scripting the transfer of administrative-console configurations and configuration property files from one PingFederate server to another.

The configcopy tool can migrate your configurations, such as from a test environment to production. It can also manage certificates for the target server.

The command-line utility, configcopy, in the <pf install>/pingfederate/bin directory, uses PingFederate's built-in Connection Management Service in conjunction with an internal Web Service to export and import connections and other configurations, and to obtain lists. For more information, see Connection Management Services.



Important:

The Connection Management Service must be activated for both the source and target servers before you can use the configcopy tool. For more information, see Configuring service authentication on page 636.



CAUTION:

For security reasons, you should disable the Connection Management Service whenever it is not in use.

Copying the key from the source to the target server

You must copy the key from the source server to the target server before you migrate data using the configcopy tool.

About this task

To copy the key from the source to the target server, you copy the needed keys from the pf.jwk file on the source server and append it to the last key in the pf.jwk file on the target server, and then restart that target server. This step only needs to be completed before the first migration.

Steps

- 1. In your PingFederate installation on the source server, open the pf.jwk file in the <pf install>/ pingfederate/server/default/data directory.
- 2. Copy the key in the file.

Make sure you copy the entire key JSON message, as in the following example.

```
{ "keys": [ { "kty": "oct", "kid": "j0PUEdAb95", "k": "AGi8Lg ewdl-
```

- 3. In your PingFederate installation on the target sever, open the pf.jwk file.
- 4. Insert a comma at the end of the last key in the file and append the source key.

For example, if the pf.jwk on the target server reads as follows.

```
{"kevs":
[{"kty":"oct", "kid":"wER9zEpaPe", "k":"i0HQr9JmsqjAX4o BQU1qGJzoLQI-
nmwp8u3GyHzTB8"}]}
```

Insert the comma and the source key as shown.

```
{"kevs":
[{"kty":"oct", "kid":"wER9zEpaPe", "k":"i0HQr9JmsqjAX4o BQU1qGJzoLQI-
nmwp8u3GyHzTB8"},
{"kty":"oct","kid":"j0PUEdAb95","k":"AGi8Lg_ewdl-
_30Cx83kDMQE9oNlhgJSa_Pc4l8JTU8"}]}
```



Note:

This is a well-formed JSON document in one line.

- 5. Save the pf.jwk file and restart the target server.
- 6. If applicable, repeat the steps above for each target PingFederate server.

Administrative console migration

Use the configcopy tool to migrate data in the administrative console.

For migrating data configured with the source server's administrative console, the configcopy tool performs these overall processing steps:

- 1. Retrieves specified connection and other configuration data (XML) from a source PingFederate server
- 2. Modifies the configuration with any changes required for the target environment, according to settings in one or more properties files, command-line arguments, or both
- 3. Imports the updated configuration into the PingFederate target server

The configcopy tool can perform these functions in real time, from server to server, or by using an intermediate file. The latter option is useful when both the source and target PingFederate servers are either not running at the same time or not accessible from the same operating system command window.



Important:

For one-time configuration transfers from one version of PingFederate to a newer version, use a complete configuration archive, either with configcopy archive export/import commands, or manually through the administrative console, or the administrative API. Other configcopy commands are not supported for this purpose.

Operational capabilities include:

- Listing of source partner connections, adapter or STS token-translator instances, outboundprovisioning channels, or datastore connections.
 - List commands include optional filter settings, when applicable.
- Copying one or more partner connections, outbound-provisioning channels, or instances of adapters or token translators.
- Copying one or more datastore connections.
- Copying server settings.
- Exporting and importing full configuration archives.

Copying configuration files

The configcopy tool supports copying configuration files containing runtime properties, including those needed for server clustering, that might have been manually customized for the source configuration and need to be migrated. The file-copy command can also copy the PingFederate internal, HSQLDB database when needed.



CAUTION:

Use the built-in HSQLDB only for trial or training environments. For testing and production environments, always use a secured external storage solution for proper functioning in a clustered environment.

Testing involving HSQLDB is not a valid test. In both testing and production, it might cause various problems due to its limitations and HSQLDB involved cases are not supported by Pingldentity.

Managing Certificates

Administrators can use the configcopy tool to perform the following certificate-management tasks on the target PingFederate server:

List source trusted certificate authorities (CAs) and target key aliases

- Copy one or all trusted CAs from the source server
- Create certificates
- Create Certificate Signing Requests (CSRs)
- Import CA-signed and PKCS-12 certificates

Using the migration tool

The migration tool, configcopy, can be used in conjunction with one or more property files to define the operational command and other parameters, including the source and target PingFederate servers, and to modify configuration settings as needed for the target environment.

About this task

Property-file templates are available for each command option in the <pf install>/pingfederate/ bin/configcopy templates directory.



See the README.txt file in the configcopy templates directory for a list of all commands and summary information. See the template files for parameters associated with each command or with use cases, as well as lists of Override Properties, which are configuration settings that can be modified in transit, where applicable.

Copies of the templates can be configured as needed and then used together, or combined into one file. Use the applicable file names as an argument when running configcopy.bat or configcopy.sh, depending on your operating system, for particular configurations, using the following command syntax.

(On Windows)

```
configcopy.bat -Dconfigcopy.conf.file=cproperties file1>;
properties file2>;...
```



Note:

When paths are included with the file names, you cannot use backslashes (\). Use forward slashes (/) or escape the backslash (\\).

(On Linux)

```
configcopy.sh -
Dconfigcopy.conf.file=cproperties file1>:cproperties file2>:...
```



Note:

The file separators are platform specific, corresponding to the syntax used for system-level path separators.

Also, you can specify any property values through command-execution arguments, using the following syntax

```
configcopy[.sh] -Dproperty>=<value> ...
```

where property> is any property named in the properties file and <value> is the value. Command-line property designations take precedence over any values set in the properties file.

Note:

Access to the Connection Management Service is password-protected. The usernames and passwords might be set in the properties file for both the source and target web services, and passwords can be obfuscated. If passwords are set in the properties file, they cannot be overridden using the command line. If a password is not set, the configcopy tool prompts for it. Usernames must always be supplied where applicable, either in the command line or in the properties file.

The configcopy utility generates its own log file, configcopy.log, located in the <pf install>/ pingfederate/log directory. You can control settings for this log, as needed, in the file configcopy.log4j2.xml, located in the bin directory.



CAUTION:

Importing connections or other discrete configurations at the target server is not subject to the same rigorous data validation performed by the administrative console during manual configuration. Although some checks are made, it is possible to create invalid connections using the connection-migration process. Therefore, you should not use the configcopy tool to create settings at the target that do not exist at the source. For connections and other configurations copied separately, the tool is designed only for modifying the values of existing source settings to make them applicable to the target environment.

To avoid errors and prevent unstable target configurations due to missing components or faulty crosscomponent references, such as invalid ID references from connection configurations to datastore configurations, adhere closely to the instructions provided in the following procedure.

Steps

- 1. Enable access to the Connection Management Service for both the source and target PingFederate servers. For more information, see Configuring service authentication on page 636.
- 2. Determine which component configurations need to be copied, including plugins.

For example, connection configurations always reference either adapter or token-translator configurations, or both, and may reference datastore configurations. These are all separate configurations, and must be copied separately in conjunction with copying connection configurations, unless they already exist at the target.

Server Settings, unless pre-configured at the target, also need to be copied over separately.

Provisioning settings can be copied separately to update target connections, as needed.

- 3. Determine whether any configuration property files or other supporting files need to be copied.
- 4. Ensure necessary plugin JAR files are installed on the target server.

The configcopy tool does not copy over these files, which include libraries for adapters, token translators, and JDBC or any custom database drivers.

The JAR files are located in either cpf install>/pingfederate/server/default/deploy or <pf install>/pingfederate/server/default/lib.

5. On the target server, ensure that signing certificates, or certificates used for XML decryption, are already in place. For more information, see Certificate and key management on page 600.

Private keys are not copied from server to server, public certificates can be copied. However, you can use configcopy to upload keys and certificates to the target server.

Look for identifying information about the target keys so you can reference the certificates in connection-copy properties.

- 6. If you have not yet installed your organization's CA-issued SSL server certificate on both the target and source servers, you can do so with a configcopy command. Otherwise, use one of the following methods to ensure that **configcopy** can contact both servers:
 - (Recommended) Install the Issuer certificate for the PingFederate SSL certificate in a separately managed trust store. Then the location of the file can be specified when running configcopy using the property configcopy.connection.trust.keystore.
 - Install the Issuer certificate for the PingFederate SSL certificate into the trust store for the Java runtime where configcopy runs.



Note:

If different SSL certificates are installed on the two servers, the configcopy tool must be able to trust both. In this case, both certificates must be installed in the trust store used by configcopy, or in the trust store for the Java runtime where configcopy runs.

7. Create properties files for the necessary commands, and for associated command-parameter values needed to copy the required configurations and any additional files.

See the README.txt file and to the properties-file templates in the <pf install>/ pingfederate/bin/configcopy templates directory.



Note:

This step and those following assume the use of properties files based on the templates provided. You can also use command-line parameters. For more information, see previous steps in this section.

8. If you are copying connections, override ID properties referencing adapter, datastores or other plugin configurations, as needed see step 2.



Important:

Ensure that the plugin configurations are either previously defined at the target or are part of the same configcopy process used to copy the connections that depend on them.

9. Create a script or run a command, or command series, that executes configcopy for each of the prepared properties files.

See the previous discussion for syntax requirements, or the README.txt file.

Outbound provisioning CLI

PingFederate provides a command-line interface (CLI) to help manage automated outbound provisioning at identity provider (IdP) sites.

Administrators can use the CLI to view the status of user provisioning, either globally or one provisioning channel at a time, and to rectify unusual situations where provisioning at the service provider (SP) might be out of sync with the enterprise user store.

The CLI tool, provmgr.bat or provmgr.sh, is located in the directory <pf install>/ pingfederate/bin. The tool interacts with the PingFederate internal datastore to maintain provisioning synchronization between the LDAP user store and the target service.

The tool creates its own log file, located at <pf install>/pingfederate/log/provmgr.log. You can control settings for this log, as needed, in the <pf install>/pingfederate/bin/ provmgr.log4j2.xml file.

The following table describes the available global and channel-specific command arguments.

Command argument	Description
Global options	
help	Describes the available options. The help also displays if you run the command with no arguments.
show-channels	Lists all channels in a table format, showing for each:
	 ID: A numeric channel ID (channel-specific commands need this ID) Name: The channel name Connection ID Status: active inactive (both the connection and the channel status are shown) User count/dirty-user-record count, such as 5000/12, which means 5000 users and 12 dirty records Source, as LDAP URL Target code
show-nodes	Shows all the provisioning-server nodes with their status and the last timestamp. Applicable only when failover provisioning is configured in the <pre><pre><pre>cpf_install</pre><pre>/pingfederate/bin/run.properties file.</pre></pre></pre>
force-node-backup	Sets the provisioner mode to FAILOVER for the associated PingFederate
Use with node number: - n <node id=""></node>	server node.
Channel-specific option	ıs
Note:	
	ecify the channel with the -c <channel-id-number> argument. For example:</channel-id-number>
provmgr -c 1show	
You can determine chann	nel ID numbers by using the global command provmgrshow-channels.
reset-group-timestamp Deletes the user-group timestamp, which forces the provisioner to provisioning group on the next cycle, even if the timestamp on that gro not actually change.	
	Depending on your LDAP server and administrative practices, you might want to schedule this command to run periodically to catch up with any users that may have been deleted, rather than deactivated, in the directory server. Some directory servers do not update the group timestamp for deleted users.
	Important:
	You should rarely need this option if users are deactivated rather than deleted. If you do need it, you might want to schedule it when other network activity is low.

Command argument	Description
reset-attribute-sync	Sets the attribute sync timestamp to 1, which forces the provisioner to look at all users for changes, not only those that have a newer timestamp on their LDAP entry.
	Important:
	This is rarely needed and might consume considerable network resources, depending on the number of users. If it is needed, you might want to schedule it when other network activity is low.
reset-values-hash	Removes the values hash for all users. The database stores a hash of attribute values for users to determine whether any values have been changed.
	This argument forces users that have a newer timestamp on their LDAP entry to be updated at the service provider, regardless of the actual field values. However, users whose recorded timestamp is unchanged are not updated.
reset-all	Equivalent to using all three of the previous arguments.
show-dirty-records	Lists all users or groups that have not been provisioned or updated at the SP site. This option is rarely needed and might consume considerable network resources, depending on the number of users. If it is needed, you might want to schedule it when other network activity is low.
show-dirty-group- records	List groups that have not been provisioned or updated at the SP site.
show-dirty-user-records	List all users that have not been provisioned or updated at the SP site.
show-group show-user	Shows all internal database fields related to the specified user or group, including transitory mapping fields, which are fields waiting to be pushed to the SP. For a user, shows all LDAP attributes retrieved from the directory server.
Use with: -u <pre>-u <pre>rovider name></pre></pre>	Note:
Or: -g <ldap guid=""></ldap>	You can obtain user or group names and GUIDs for dirty records, as needed, using any of theshow-dirty-* options, described above.
	The LDAP GUID, if used and if it is binary, should be entered in hexadecimal format, as shown in log files.
	provmgr.shshow-user -u user@example.com provmgr.shshow-user -g ffd448643f812b43a0bee2504173f0
clear-dirty-records	Clears the dirty flag on all records.
clear-dirty-group- records	Clears the dirty flag on all group records.
clear-dirty-user-records	Clears the dirty flag on all user records.
delete-dirty-records	Removes all dirty records from the internal store.
delete-dirty-group- records	Removes all dirty group records from the internal store.

Customizable user-facing pages

PingFederate supplies HTML templates, located in the $<\!pf_install>$ /pingfederate/server/default/conf/template directory, to provide information to the end-users or to request user input when processing their requests.

The PingFederate HTML templates utilize the Velocity template engine, an open-source Apache project. For more information about Velocity, please refer to the Velocity project documentation on the Apache website at https://velocity.apache.org.

You can modify most of these pages in a text editor to suit the particular branding and informational needs of your PingFederate installation. CSS and images for these pages are included in the template/assets subdirectory. Each page contains both Velocity constructs and standard HTML. The Velocity engine interprets the commands embedded in the template page before the HTML is rendered in the user's browser. At runtime, PingFederate supplies values for the Velocity variables used in the template.

Each template contains specific variables that can be used for rendering the associated web page. You can see the variables and usage examples in the comments of each template. The following table describes variables that are available across all templates.

Variable	Description
\$escape	A utility class that can be used to escape String variables inserted into the template, such as \$escape.escape(\$client.name) where \$client.name is one of the variables available in the oauth.approval.page.template.html template file.
	Use <code>\$escape.forJavaScript(\$variable)</code> when passing String variables into a JavaScript code block or an event handler within a template, such as <code>window.location.replace("\$escape.forJavaScript(\$wreply)")</code> as seen in the <code>sourceid-wsfed-idp-signout-cleanup-template.html</code> template file.
	Important:
	Use the \$escape variable to escape external data, such as request parameters, to mitigate the risk of potential cross-site scripting (XSS) attacks.
\$HttpServletRequest	A Java object instance of <code>javax.servlet.http.HttpServletRequest</code> . Used to add additional knowledge about the request that is otherwise unavailable in the template, such as the <code>User-Agent</code> HTTP header.
\$HttpServletResponse	A Java object instance of <code>javax.servlet.http.HttpServletResponse</code> . Used to modify the response in the template, such as setting additional browser cookies.
\$locale	A Java object instance of <code>java.util.Locale</code> that represents a user's country and language. Used to customize the end-user experience. For example, the locale is used to display content in the user's preferred language.
\$CurrentPingFedBaseU	RLThe host name found in the request, provided that it matches either the PingFederate's base URL or one of the configured virtual host names.
\$PingFedBaseURL	The PingFederate base URL.
	For most deployments, use the \$CurrentPingFedBaseURL variable instead of the \$PingFedBaseURL variable.
\$templateMessages	Used to localize messages in the template, based on user's Locale, an instance of com.pingidentity.sdk.locale.LanguagePackMessages. For more information, see the Javadoc for the LanguagePackMessages class in the directory <pre>cf_install>/pingfederate/sdk/doc.</pre>
\$TrackingId	The user's session tracking ID.



Important:

Changing Velocity or JavaScript code is not recommended.

At runtime, the user's browser is directed to the appropriate page, depending on the operation being performed and where the related condition occurs. For example, if a single sign-on (SSO) error occurs during identity provider (IdP)-initiated SSO, the user's browser is directed to the IdP's SSO error-handling page.

Applications can override the PingFederate server-hosted pages provided specifically for SSO and single logout (SLO) errors by specifying a URL value in the relevant application endpoint's InErrorResource parameter. Administrators can override SSO and SLO success pages by specifying default URLs on the SP Default URLs window (Applications # Integration # SP Default URLs) or the IdP Default URL window (Authentication # Integration # IdP Default URL).



If you have a clustered PingFederate environment, copy the customized, and localized, templates to each

IdP user-facing pages

PingFederate has a variety of customizable user-facing page templates that apply to identity provider (IdP) pages. The templates are organized by HTML Form Adapter, Kerberos Adapter, single sign-on (SSO), single logout (SLO), WS-Federation, and OpenID Connect (OIDC).

HTML Form Adapter

Page title and template file name	Purpose	Туре	Action
Sign On or Choose an Account identifier.first.template.html	Prompts a user to provide their username when an Identifier First Adapter instance is invoked to handle a sign on request.	Normal	User input required
Sign On html.form.login.template.html	Displays a customizable user sign-on form when an HTML Form Adapter instance is invoked to handle a sign on request.	Normal	User input required
	If the invoked HTML Form Adapter instance is associated with a local identity profile that is configured to support authentication using third- party identity providers, those identity providers are displayed on the sign-on page as well.		
Change Password html.form.change.password.template.ht	Displayed when a user attempts to change their mi password through the HTML Form Adapter.	Normal	User input required
Change Password html.form.message.template.html	Displayed when a user successfully changes their password.	Normal	User input required
Password Expiring html.form.password.expiring.notification	Displayed to warn an authenticated user that the template html password associated with the account is about to expire.	Normal	User input required
Password Management System Message html.form.message.template.html	Displayed when a user is being redirected to a password management system to change their password.	Normal	User input required

Page title and template file name	Purpose	Туре	Action
Account Recovery	Displayed when a user	Normal	User input
forgot-password.html	attempts to reset their password through the HTML Form Adapter.		required
	If the user enters a username in the sign-on form, the username carries over to this form. Otherwise, the user must enter their username to begin the self-service password reset process.		
Account Recovery forgot-password-resume.html	Displayed to prompt a user to enter the one-time password sent through a notification or to notify a user to refer to the notification for password reset instructions.	Normal	User input required
	This template is applicable when the password reset type is Email One-Time Link, Email One-Time Password, or Text Message for the invoked HTML Form Adapter instance.		
Reset Your Password forgot-password-change.html	Displayed to prompt a user to define a new password.	Normal	User input required
Account Recovery forgot-password-success.html	Displayed when a user successfully resets their password.	Normal	User input required
Account Recovery forgot-password-error.html	Displayed when a password reset attempt fails.	Error	None
Unlock Your Account account-unlock.html	Displayed when a user successfully unlocks their account through the HTML Form Adapter.	Normal	User input required
	This page also prompts the user to retain the current password, or reset it.		
Security Question html.form.login.challenge.template.html	Displays a configurable challenge form for two- step authentication. For example, this template can be used to create a RADIUS challenge form when using the RADIUS Username/Password Credential Validator.	Normal	User input required

Page title and template file name	Purpose	Туре	Action
User Consent consent-form-template.html	Displayed when a request requires a user's consent for an SSO to an SP.	Normal	User input required
Logout Confirmation idp.slo.confirm.page.template.html	Displayed when a user initiates a logout request.	Normal	User input required
nap.sio.somm.page.template.mim	Applicable only if such confirmation is required, as configured on the Authentication # Integration # IdP Default URL window.		
Sign Off	Displayed when a user	Normal	None
idp.logout.success.page.template.html	successfully signs off in a configuration where the Logout Path field is configured but the Logout Redirect field is not.		
Create Your Account	Displayed when a user	Normal	User input
local.identity.registration.html	requests to register for a local account.		required
	Applicable only if the invoked HTML Form Adapter instance is associated with a local identity profile that is configured to support self-service registration.		
Manage Your Profile	Displayed when an	Normal	User input
local.identity.profile.html	authenticated user accesses the profile management endpoint.		required
	Applicable only if the invoked HTML Form Adapter instance is associated with a local identity profile that is configured to support self-service profile management.		

Page title and template file name	Purpose	Туре	Action
Email Verification local.identity.email.verification.sent.html	Displays a notification that an email ownership verification message has been sent when an authenticated user accesses the email ownership verification endpoint.	Normal	None
	Applicable only if the invoked HTML Form Adapter instance is associated with a local identity profile that is configured to offer users the opportunity to verify the ownership of the email address associated with the accounts.		
Email Verified local.identity.email.verification.success.h	Displays a confirmation that the user has successfully verified the ownership of the email address associated with the account.	Normal	None
	Applicable only if the invoked HTML Form Adapter instance is associated with a local identity profile that is configured to offer users the opportunity to verify the ownership of the email address associated with the accounts.		
Email Verification Error local.identity.email.verification.error.html	Displays that the user failed to verify the ownership of the email address associated with the account.	Error	User can request another verification
	Applicable only if the invoked HTML Form Adapter instance is associated with a local identity profile that is configured to offer users the opportunity to verify the ownership of the email address associated with the accounts.		email by accessing the email ownership verification endpoint or the profile management page (if enabled). Authentication is required.
			Alternatively, the user can contact their IT administrators for further assistance.

Page title and template file name	Purpose	Туре	Action
Username Recovery username.recovery.template.html	Displays to prompt the user to enter an email address to recover the username associated with the account.	Normal	User input required
	Applicable only if the invoked HTML Form Adapter instance is configured to support self-service username recovery.		
Username Recovery username.recovery.info.template.html	Displays to notify the user to retrieve the notification message with the recovered username.	Normal	User should retrieve the notification message with
	Applicable only if the invoked HTML Form Adapter instance is configured to support self-service username recovery.		the recovered username.

Kerberos Adapter

Page title and template file name	Purpose	Туре	Action
Error kerberos.error.template.html	Displays an error page to provide standardized information to the end user when the authentication attempt fails.	Error	Consult log
(No title) meta.refresh.template.html	Facilitates the failover mechanism from a Kerberos Adapter instance to the next phase when it is part of a Composite Adapter instance configuration or an authentication policy.	Normal	None

Single sign-on and logout

Page title and template file name	Purpose	Туре	Action
Select Authentication System	Displayed when multiple	Normal	User input
sourceid-choose-idp-adapter-form-template.html	authentication sources are applicable and no preference is submitted as part of the request.		required
Sign On Error	Displayed when IdP-initiated or	Error	Consult log
idp.sso.error.page.template.html	adapter-to-adapter SSO fails and no other SSO error landing page is specified.		and web developer
Sign Off Successful	Displayed when an SLO	Normal	None
idp.slo.success.page.template.html	request succeeds and no other SLO success landing page is specified.		

Page title and template file name	Purpose	Туре	Action
Sign Off Error idp.slo.error.page.template.html	Displayed when an SLO request fails and no other SLO error landing page is specified.	Error	User should close the browser

WS-Federation and OpenID Connect

Page title and template file name	Purpose	Туре	Action
Working sourceid-wsfed-http-post-template.html	Used to auto-submit a WS-Federation assertion to the SP. If JavaScript is disabled, the user is prompted to click a button to POST the assertion directly.	Normal	None
	This page is normally not displayed if JavaScript executes properly.		
Signing off sourceid-wsfed-idp-signout-cleanup-invisible-template.html	WS-Federation and OIDC client IdP sign-out processing page. No HTML is rendered in the browser.	Normal	None
Sign Off Successful sourceid-wsfed-idp-signout-cleanup-template.html	Indicates user signed out of the IdP under the WS-Federation protocol and lists each successful SP logout, when applicable.	Normal	None
	Also displays when an OIDC client sends a logout request to the /idp/startSLO.ping endpoint to initiate an Asynchronous Front-Channel Logout process.		

SP user-facing pages

PingFederate has a variety of customizable user-facing page templates that apply to service provider (SP) pages. Each template contains specific variables that can be used for rendering the associated web page.

Account linking

Page title and template file name	Purpose	Туре	Action
Link Your Account LocalIdPasswordLookup.form.template.h	Used to authenticate a user at the SP when an account link needs to be established.	Normal	None
Account Unlinked TerminateAccountLinks.page.template.h	Communicates a user's successful defederation operation.	Normal	None

Page title and template file name	Purpose	Туре	Action
Select Identity Provider sourceid-saml2-idp-selection- template.html	The user requested SP-initiated single sign-on (SSO), but the identity provider (IdP) partner was not specified in the appropriate query parameter or cookie. This page allows the user to select the IdP manually. Based on the user's selection, the server redirects the browser to the appropriate IdP partner's SSO service.	Normal	User must make selection
Please Specify Target sp.sso.success.page.template.html	Displayed when an SSO request succeeds but no target-resource parameter is specified by the incoming URL, and no default URL is set on the Applications # Integration # SP Default URLs window.	Error	Consult web developer or specify default URL
Sign On Error sp.sso.error.page.template.html	Displayed when SP-initiated SSO fails, or IdP-initiated SSO fails on the SP side, and no other SSO error landing page is specified.	Error	Consult log and web developer
Sign Off Successful sp.slo.success.page.template.html	Displayed when a single logout (SLO) request succeeds and no other SLO success landing page is specified.	Normal	None
Sign Off Error sp.slo.error.page.template.html	Displayed when an SLO request fails and no other SLO error landing page is specified.	Error	User should close the browser

WS-Federation

Page title and template file name	Purpose	Туре	Action
Signed Off	Displays the user's sign-out	Normal	None
sourceid-wsfed-sp-signout-cleanup- template.html	status.		
Unable to Authenticate	Displayed when an	Error	Consult log
sourceid-wsfed-idp-exception- template.html	authentication challenge fails during WS-Federation processing.		and web developer

PingFederate allows customizable user-facing windows that can be applied to both the identity provider (IdP) and service provider (SP) pages.

Page title and template file name	Purpose	Туре	Action
Sign On AbstractPasswordIdpAuthnAdapter.form	Challenges user for credentials when authentication can template html take place via HTTP Basic authentication or an HTML form, depending on the operational mode.	Normal	User must sign on
Submit Form form.autopost.template.html	Whenever the server posts a form, this template is used to auto-submit the form. If JavaScript is disabled, the user is prompted to click a button to post the form manually.	Normal	None
	This page is normally not displayed if JavaScript executes properly.		
Multiple Sign-On Delay speed.bump.template.html	Displayed to indicate that simultaneous single sign-on requests from multiple browser tabs are in progress.	Normal	User can switch to the browser tab that is actively waiting for the user to complete the authentication requirement or resubmit the request.
Error general.error.page.template.html	Indicates that an unknown error has occurred and provides a error reference number and, optionally, an error message.	Error	Consult log Contact Ping Identity Support if unresolved
Error generic.error.msg.page.template.html	General error, with error code.	Error	Consult log and check configuration Contact Ping Identity Support if
Error http.error.page.template.html	Indicates that an HTTP error has occurred and provides the HTTP status code.	Error	unresolved Consult log Contact Ping Identity Support if unresolved

OAuth user-facing pages

The PingFederate OAuth authorization server provides five windows that are presented to end-users, or resource owners, during certain OAuth transactions. You can customize and brand these windows as needed..

Page title and template file name	Purpose	Message type	Action
Client Access oauth.access.grants.page.template.html	Provides a means for the end users, or resource owners, to revoke persistent access grants.	Normal	User input required

Customizable email notifications

PingFederate delivers messages to administrators and end-users based on customizable notification publisher settings.

Each component that is capable of triggering or handling events can use a different notification publisher instance to deliver its messages. For example, you can select an SMTP Notification Publisher instance to deliver messages to your end users in an HTML Form Adapter instance and another SMTP Notification Publisher instance to deliver licensing messages to your fellow administrators.

When a component is configured to deliver its messages based on an SMTP Notification Publisher instance configuration, PingFederate creates notification messages based on template files located in the <pf install>/pingfederate/server/default/conf/template/mail-notifications directory. Each template file is a combination of variables and HTML codes. You can modify these template files in a text editor to suit the particular branding requirements, as needed.



Note:

If you have a clustered PingFederate environment, copy the customized templates to each node.



You can also configure a component to use an Amazon SNS Notification Publisher instance to deliver notification messages. If so, refer to Configuring an Amazon SNS Notification Publisher instance on page 927 and Event types and variables on page 927 for more information about this messaging model and the handling of notification messages.

Local administrative account management events

PingFederate generates a variety of email notifications for local administrative account management events.

Notification for account management events is configurable from System # Serverin the Administrative Accountswindow.



Account management events are only applicable when native authentication is enabled for the administrative console, the administrative API, or both in the <pf install>/pingfederate/bin/ run.properties file. If you are using an alternative console authentication, notifications, if any, such as password changes, are handled by the third-party system.

Email subject and template file name	Event	Action
PingFederate Notification Settings Change Notification message-template-notifications.html	An administrator has turned off the Notify Administrator of Account Changes option.	Ensure this change is approved and legitimate.
message-template-notifications.html	PingFederate generates a notification message to all administrators.	
	The message includes the username of the administrator who made the change.	
PingFederate Email Change Notification	An administrator's email	Ensure this change is
message-template-email.html	address has been updated by another administrator.	approved and legitimate.
	PingFederate generates a notification message to the previous email address and another notification to the new email address.	
	The message includes the username of the administrator who made the change.	
PingFederate Password Change Notification	An administrator's password has been changed.	Ensure this change is approved and legitimate.
message-template-password.html	PingFederate generates a notification to the administrator whose password has been changed.	
	The message includes the username of the administrator who made the change.	

Certificate events

PingFederate sends email notifications for the creation, update, and expiration of certificates.

PingFederate also sends email notifications for the creation and activation of new pending certificates when automatic rotation for self-signed certificates is enabled in the Security # Certificate & Key

Email subject and template file name	Event	Action
A PingFederate Certificate Is About to Expire message-template-cert-warning.html	re PingFederate generates a	Create a new certificate and work with the applicable partners to update the expiring certificate.
	Notifications window. The message includes the details of the certificate and the connections associated with it.	If a self-signed certificate is used for signing or decryption, consider enabling certificate rotation while creating the new certificate.
		For a clustered PingFederate environment, replicate the configuration using the administrative console.
A PingFederate Certificate Has Expired message-template-cert-expire.html	A certificate expired. PingFederate generates a notification upon the expiration of a certificate.	Create a new certificate and work with the applicable partners to update the expiring certificate.
	The message includes the details and the connections associated with the certificate.	If a self-signed certificate is used for signing or decryption, consider enabling certificate rotation while creating the new certificate.
		For a clustered PingFederate environment, replicate the configuration using the administrative console.
A New PingFederate Certificate Has Been Created message-template-cert-rotation.html	A new pending certificate has been created for signing or decryption.	Work with the applicable partners to update the expiring certificate.
33.1.013.1.013.1.013.1.013.1.013.1.013.1.013.1.013.1.013.1.013.1.013.1.013.1.013.1.013.1.013.1.013.1.013.1.013	PingFederate generates a notification when a new pending certificate is created.	PingFederate supports providing metadata for Browser SSO connections.
	The message includes the details of the current certificate, the details of the new certificate, the activation date, and the connections that will be affected when the new certificate is activated.	For a clustered PingFederate environment, replicate the configuration using the administrative console.

Email subject and template file name	Event	Action
A PingFederate Certificate Has Been Updated	A new certificate for signing or decryption is activated.	None, unless the applicable partners have not been
message-template-cert-deactivation.html	PingFederate generates a notification when the new certificate is activated.	notified or configuration has not been replicated in a clustered PingFederate environment.
	The message includes the details of the new certificate and the affected connections.	

SAML metadata update events

When notification for SAML metadata update events is enabled in the **Runtime Notifications** window, PingFederate sends the notifications to the email address that has been configured for the event.

PingFederate supports automatic reloading of SAML metadata, which streamlines the maintenance of SAML connections.

Email subject and template file name	Event	Action
Your \${CONNECTION_NAME} \${SP_IDP} Connection in PingFederate Has Been Updated message-template-metadata-	PingFederate has updated a connection based on the partner's SAML metadata.	For a clustered PingFederate environment, replicate the configuration using the administrative console.
updated.html		
Please Review Your Updated \${CONNECTION_NAME} \${SP_IDP} Connection in PingFederate	a connection based on the partner's SAML metadata. However, some settings are out of sync.	Review the settings that are out of sync and make changes as needed.
message-template-metadata-updated- out-of-sync.html		For a clustered PingFederate environment, replicate the configuration using the administrative console.
Your \${CONNECTION_NAME} \${SP_IDP} Connection in PingFederate Is Out of Sync	A connection is out of sync with the changes found in the partner's SAML metadata.	Review the settings that are out of sync and make changes as needed.
message-template-metadata-out-of- sync.html		For a clustered PingFederate environment, replicate the configuration using the administrative console.
Your Metadata Couldn't Be Downloaded from \${METADATA_URL_NAME}	PingFederate failed to download the SAML metadata	Consult log. Verify and update the metadata URL and its
message-template-metadata-url- notification.html	•	corresponding verification certificate.
Your Metadata for PingFederate \${SP_IDP} Connection \${CONNECTION_NAME} Wasn't Found	The partner's metadata URL did not return the expected metadata for a given	Consult log. Verify and update the metadata URL.
message-template-metadata-url-entity-id-missing.html	connection.	

There is a list of different licensing events and different actions you can take to resolve each event.

When you enable notifications for licensing events, PingFederate sends license expiry information to the recipient configured for the event on the Runtime Notifications window.



To check the details of your license, sign on to the administrative console, navigate to the user icon in the upper-right corner of the administrative console, and click About the list. The license summary displays in a pop-up browser window.



Important:

If the license specifies an expiration date, the license expires at the beginning of that day.

Email subject and template file name	Event	Action
Your PingFederate License Is About to Expire	The current license is about to expire.	sales@pingidentity.com to
message-template-warn.html	The email notification is sent 60 days ahead of the expiration.	renew the license before the expiration.
	Applicable to licenses without connection groups.	
Your PingFederate License Is About to Expire message-template-group-warn.html	One of the connection groups in the current license is about to expire.	Contact sales@pingidentity.com to renew the license before the
message-template-group-warn.ntm	The email notification is sent 60 days ahead of the expiration.	expiration.
	Applicable to licenses with connection groups.	
PingFederate License Expiration	The current license expired.	Contact
Notification message-template-grace.html	The email notification is sent upon the expiration.	sales@pingidentity.com to renew the license.
	Applicable to licenses without connection groups.	
Your PingFederate License Has Expired message-template-group-grace.html	One of the connection groups in the current license expired.	Contact sales@pingidentity.com to
The state of the s	The email notification is sent upon the expiration.	renew the license.
	Applicable to licenses with connection groups.	

HTML Form Adapter events

The HTML Form Adapter offers self-service account management tools for password management, account recovery, username recovery, and email ownership verification.

When you configure an HTML Form Adapter instance to deliver its messages based on an SMTP Notification Publisher instance configuration, administrators can optionally customize and localize the following template files with branding controls to provide a consistent brand experience to end users across multiple user populations.

Email subject and template file name	Event	Action	
Password Change Notification message-template-end-user-password- change.html	A user's password has been changed successfully through an instance of the HTML Form Adapter.	Users should contact their IT administrators if they have not made any attempts to change their password.	
Password Reset message-template-forgot-password-link.html	A user has initiated a self- service password reset request.	Users should contact their IT administrators if they have not made any attempts to reset their password.	
	Applicable when the password reset type for the invoked HTML Form Adapter instance is set to Email One-Time Link .	anon pacemera.	
Password Reset message-template-forgot-password- code.html	A user has initiated a self- service password reset request.	Users should contact their IT administrators if they have not made any attempts to reset	
	Applicable when the password reset type for the invoked HTML Form Adapter instance is set to Email One-Time Password or Text Message.	their password.	
Password Reset Completed	A user's password has been	Users should contact their IT	
message-template-forgot-password-complete.html	reset successfully through an instance of the HTML Form Adapter.	administrators if they have not made any attempts to reset their password.	
	Applicable when the password reset type for the invoked HTML Form Adapter instance is set to any method except None .		

Email subject and template file name	Event	Action
Username Recovery message-template-username- recovery.html	A user has initiated a self-service username recovery request and provided an email address through an instance of the HTML Form Adapter. If PingFederate can locate the user record using such email address and other requirements are met, PingFederate uses this template to generate an email message containing the recovered username and sends it to the user at the email address provided by the user.	Users should contact their IT administrators if they have not made any attempts to recover the username associated with their accounts.
	Applicable only if the invoked HTML Form Adapter instance is configured to support self-service username recovery.	

Customizable text message

You can customize text messages in PingFederate for a unique experience.

If you have configured your self-service password reset and, optionally, your self-service account unlock to use the text message option, you can customize and localize the text message for a unique experience. The default message is stored as a property in the <pf install>/pingfederate/server/default/ conf/language-packs/pingfederate-sms-messages.properties file.



Note:

If you have a clustered PingFederate environment, copy the customized, and localized, templates to each node.

Localizing messages for end users

PingFederate supports localization for several types of messages for end users.

About this task

Administrators can localize the following message types:

- on-screen messages
- · email messages
- text messages (SMS)
- authentication API error messages

The English contents for each message type are in the <pf install>/pingfederate/server/ default/conf/language-packs directory.

Message type	Default message file
On-screen messages	pingfederate-messages.properties
Email messages	pingfederate-email-messages.properties

Message type	Default message file
Text messages (SMS)	pingfederate-sms-messages.properties
Authentication API error messages	authn-api-messages.properties

Steps

- 1. Create a copy of the associated default message file in the language-packs directory.
- 2. Provide translated text in place of English and then append the standard language tag to the base file name, as indicated in the browser settings.

For example, to localize the user-facing window messages in French, rename the translated copy of the localization file to pingfederate-messages fr.properties.

If the system language of the PingFederate server is not English, copy the default (English) message file and append en at the end of the file name for any English users, such as from pingfederatemessages.properties to pingfederate-messages en.properties, translate the default message file for the local users, and provide additional translations as needed.

3. If you want to include a region, append the capitalized abbreviation to the standard language tag with an underscore between them.

For example, to localize the text messages in Canadian French, rename the translated copy to pingfederate-sms-messages fr CA.properties.



Note:

The capitalization and underscore usage might not correspond to the way regions are listed in browser settings. However, the usage is required by the Java-based localization implementation.

4. If you have a clustered PingFederate environment, copy the localized message files to each node.

Result

For on-screen and email messages, developers can also customize the look and feel of the templates by using localization variables in logic statements to control fonts, color, and other style elements. For more information, see the template files for examples.



To maximize performance, PingFederate caches localized UI strings on start-up. For testing new localization implementations, an administrator can temporarily turn off caching by changing the value of the cache-language-pack-messages element to false in the <pf install>/pingfederate/ server/default/data/config-store/locale-options.xml file.

Return the value to true when testing is complete. You must restart the server after changes to configuration files.

Locale overrides by cookies

An administrator or web developer might want to provide end-users a means of overriding browser language preferences temporarily by setting cookies, such as creating a company web portal link for users to click instead of manually changing their browser options.

By default, the PingFederate localization framework supports overriding the locale using a cookie named pf-accept-language. The cookie value must conform to the guidelines defined under IETF BCP 47. For more information about the Java core method that PingFederate uses to parse the cookie, see Locale.forLanguageTag(String languageTag).

This locale-override behavior is the default implementation of the Java interface, LocaleOverrideService, defined in the PingFederate SDK. For more information, see the Javadoc for that interface in the <pf install>/pingfederate/sdk/doc directory.

PingFederate displays the language indicated in the cookie if the language is supported in the languagepacks directory. If the matching localization file is not found, PingFederate defaults to the browser settings.

Retrieval of localized messages

Retrieval of localized messages is supported through the LanguagePackMessages class available in the PingFederate SDK. An instance of this class is passed into every template file and made available for use there. For more information, see the Javadoc for the class in the <pf install>/pingfederate/sdk/ doc directory.

Configuring a password policy

PingFederate applies a configurable policy to passwords, pass phrases, and shared secrets defined by administrators in the administrative console.

About this task

These fields include, but are not limited to:

- Passwords used by HTTP Basic authentication for:
 - Inbound SOAP messages from partners via back-channel calls
 - WS-Trust STS
- Shared secrets used by the credentials defined for:
 - Attribute Query
 - Java Management Extensions (JMX)
 - Connection Management
 - Single sign-on (SSO) Directory Service
- Passwords used by instances of the Simple Username Password Credential Validator (PCV)
- Passwords used for encrypting certificates exported with their private keys
- Pass phrases used by identity provider (IdP) Discovery
- Passwords used by administrative console credentials when native authentication is used



Note:

Passwords external to PingFederate, such as passwords used by instances of the datastores, are not subject to this password policy.

Steps

- 1. Edit the <pf install>/pingfederate/server/default/data/config-store/passwordrules.xml file.
- 2. Save the changes.
- 3. Restart PingFederate.

For a clustered PingFederate environment, perform these steps on the console node. You do not have to change or restart PingFederate on the engine nodes.

Managing cipher suites

You can enable, disable, and re-order cipher suites in PingFederate.

About this task

The SSL/TLS server-client handshake involves negotiating cipher suites to use for encryption and decryption on each side of a secured transaction. You can find cipher suites in the following configuration files:

- com.pingidentity.crypto.SunJCEManager.xml
- com.pingidentity.crypto.AWSCloudHSMJCEManager.xml
- com.pingidentity.crypto.LunaJCEManager.xml
- com.pingidentity.crypto.NcipherJCEManager.xml
- com.pingidentity.crypto.BCFIPSJCEManager.xml

These cipher-suite configuration files are located in the <pf install>/server/default/data/ config-store directory. These files comment out weaker cipher suites. To ensure the most secure transactions, retain this cipher-suite configuration.



Important:

Due to the import restrictions of some countries, Oracle Server Java SE Runtime Environment (JRE) 8 has built-in restrictions on available cryptographic strength (key size). To use larger key sizes, enable the Java Cryptography Extension (JCE) unlimited strength jurisdiction policy. For more information, see the Java 8 release notes in Oracle's documentation.

For Oracle Java SE Development Kit 11, the JCE jurisdiction policy defaults to unlimited strength. For more information, see the Oracle JDK Migration Guide in Oracle's documentation.

Starting with PingFederate 9.1, cipher suites are selected based on the order that they are listed in the cipher-suite configuration file for new installations. For upgrades, you can enable the same selection mechanism as well.

Steps

- Choose one of the following actions. Choose from:
 - Edit the applicable cipher-suite configuration file.
 - Save your changes.
 - Restart PingFederate.
 - a. Edit the applicable cipher-suite configuration file.
 - b. Save your changes.
 - c. Restart PingFederate.

For a clustered PingFederate environment, perform these steps on the console node, and then click Replicate Configuration on System # Server # Cluster Management.



Important:

For each engine node, restart PingFederate to load the changes made in the cipher-suite configuration file after the configuration is replicated.

- To enable cipher-suite selection based on listing order after an upgrade, follow these steps.
 - a. Create a new text file with the following content.

```
<?xml version="1.0" encoding="UTF-8"?>
<c:config xmlns:c="http://www.sourceid.org/2004/05/config">
    <c:item name="prefer-server-cipher-suites">true</c:item>
</c:config>
```

- b. Save this file as cipher-suite-settings.xml in the <pf install>/pingfederate/ server/default/data/config-store directory.
- c. Restart PingFederate.

For a clustered PingFederate environment, perform these steps on the console node, and then click Replicate Configuration on System # Server # Cluster Management.



Important:

For each engine node, restart PingFederate to load the changes made in the cipher-suitesettings.xml file after the configuration is replicated.

Manage externally stored authentication sessions

Authentication sessions control when previously authenticated users are redirected back to the authentication sources on subsequent requests for browser-based single sign-on (SSO) and PingFederate user-facing applications.

When you enable authentication sessions, PingFederate maintains session data in memory. PingFederate also supports maintaining session data both in memory and on an external storage. This optional capability allows your organization to support use cases where a longer session duration or a greater resilience against restarts of PingFederate and browsers is desired.

PingFederate supports storing persistent authentication sessions on a database server or a PingDirectory server. When stored on a database server, the default cleanup task removes expired authentication sessions once a day. If stored on a PingDirectory server, configure a cleanup plugin in PingDirectory to suit the needs of your organization.

Managing authentication sessions stored in the database

PingFederate uses a cleanup task to remove expired authentication sessions from the configured database once a day. The cleanup task determines whether a session can be removed by looking at the session's expiration timestamp and the current time.

About this task

Any session that has an expiration timestamp older than the current time by a configurable offset is subject to removal. As needed, the cleanup task can look at the session's last activity timestamp instead. The cleanup task removes 500 expired sessions at a time until all expired sessions are removed. If expired sessions are growing rapidly, you can optionally increase the frequency of the cleanup task.



Increasing the frequency of the cleanup task or the number of expired sessions to be removed per batch (or both) adds more workload to your storage server. Make changes gradually to observe the impact.



Important:

In a clustered PingFederate environment, the cleanup task runs only on the console node. If adjustments are required, make them on the console node. No changes are required on any of the engine nodes.

Steps

- 1. Optional: Adjust the frequency of the cleanup task.
 - a. Edit the <pf_install>/pingfederate/server/default/data/config-store/timer-intervals.xml file.
 - b. Update the StoredSessionCleanerInterval value, in milliseconds.
 - The default value is 86400000, which is 24 hours.
 - c. Save your changes.
- 2. Optional: Configure other cleanup options.
 - a. Edit the <pf_install>/pingfederate/server/default/data/config-store/ org.sourceid.saml20.service.session.data.impl.SessionStorageManagerJdbcImpl.xml file.

See the following table for more information about each field.

Field	Description
ExpiredSessionGroupBa TcheSinze mber of expired authentication sessions to be removed per batch.	
	The default value is 500.

Field

Description

ExpirationTimeColumnNamecolumn of which its value determines whether an authentication session has expired in the context of the cleanup task. Valid options are expiry time and last activity time.

expiry_time

Set to <code>expiry_time</code> if the cleanup task should only remove persistent authentication sessions that have expired.

The cleanup task determines if a session can be removed by looking at the session's expiration timestamp and the current time. If the expiration timestamp is older than the current time by the number of minutes specified by the <code>ExpirationTimeOffsetMins</code> field, the session is subject to removal.

last_activity_time

Set to $last_activity_time$ if the clean task should remove persistent authentication sessions that have been left idle.

The cleanup task determines if a session can be removed by looking at the session's last activity timestamp and the current time. If the last activity timestamp is older than the current time by the number of minutes specified by the <code>ExpirationTimeOffsetMins</code> field, the session is subject to removal.

For example, if PingFederate should remove persistent authentication sessions for which the last activity time is more than three weeks ago, set the ExpirationTimeColumnName value to last_activity_time and the ExpirationTimeOffsetMins value to 30240.

The default value is expiry_time.

ExpirationTimeOffsetMin\(\bar{\star}\) he offset, in minutes, relative to the current time.

The default value is 10.

- b. Save your changes.
- 3. If you have made any changes, restart PingFederate.

In a clustered PingFederate environment, you do not have to change or restart PingFederate on any of the engine nodes.

When storing persistent authentication sessions on a PingDirectory server, you must also configure a cleanup plugin in PingDirectory to remove expired authentication sessions from your directory server.

Steps

1. Disable the PingFederate cleanup task.



Important:

For a clustered PingFederate environment, make these changes on the console node. None of the engine nodes require any changes.

- a. Edit the <pf install>/pingfederate/server/default/data/config-store/timerintervals.xml file.
- b. Update the StoredSessionCleanerInterval value to 0.
- c. Save your changes.
- d. Restart PingFederate.
- 2. Sign on to the PingDirectory administrative console.
- 3. Go to Configuration # Plugin Root.
- 4. On the Plugin Root window, click New Plugin, and then select Purge Expired Data Plugin.
- 5. Configure a new instance of the **Purge Expired Data Plugin**.

See the following table for information about each required field.

Field	Description
Name	The name of this plugin instance.
Enabled	The status of this plugin instance.
	Select the check box to enable this plugin instance. Clear the check box to disable this plugin instance.
	This check box is not selected by default.

Field	Description
Datetime Attribute	The attribute value determines whether an authentication session has expired in the context of this plugin instance. Valid options are pf-authn-session-group-expiry-time and pf-authn-session-group-last-activity-time.
	pf-authn-session-group-expiry-time
	Set to pf-authn-session-group-expiry-time if this plugin instance should only remove persistent authentication sessions that have expired.
	This plugin instance determines if a session can be removed by looking at the session's expiration timestamp and the current time. If the expiration timestamp is older than the current time by the number of minutes specified by the Expiration Offset field, the session is subject to removal.
	pf-authn-session-group-last-activity-time
	Set to pf-authn-session-group-last-activity-time if the clean task should remove persistent authentication sessions that have been left idle.
	This plugin instance determines if a session can be removed by looking at the session's last activity timestamp and the current time. If the last activity timestamp is older than the current time by the number of minutes specified by the Expiration Offset field, the session is subject to removal.
	For example, if PingFederate should remove persistent authentication sessions for which the last activity time is more than three weeks ago, set the Datetime Attribute value to pf-authn-session-group-last-activity-time and the Expiration Offset value to 3 w.
Datetime Format	The format of the attribute specified in the Datetime Attribute field.
	Select generalized-time from the list.
	The default selection is generalized-time .
Expiration Offset	The offset relative to the current time.
	Enter an integer to indicate the time value, followed by its unit of measurement.
	This field has no default value.
Purge Behavior	The method how this plugin instance removes expired data.
	Select subtree-delete-entries from the list.
	This field has no default selection.
Polling Interval	The frequency of which this plugin instance should be run.
	Enter an integer to indicate the time value, followed by its unit of measurement.
	This field has no default value.

Click Save.

OAuth persistent grants cleanup

PingFederate provides two cleanup tasks for persistent grants. One task manages expired grants, while another task caps the number of grants based on a combination of user, client, grant type, and authentication context.

Persistent authorizations include those obtained by OAuth clients in the following ways:

 Grants obtained or updated using the authorization code, resource owner credentials, or device authorization grant type, in conjunction with the refresh token grant type



Note:

If the use cases involve mapping attributes from authentication sources, such as IdP adapter instances or IdP connections, or password credential validator (PCV) instances to the access tokens, directly or through persistent grant-extended attributes, storing these attributes from authentication sources and their values along with the persistent grants maintains them for reuse when clients subsequently present refresh tokens for new access tokens.

Grants obtained or updated by using the implicit grant type, for which PingFederate is configured to reuse existing persistent grants



Note:

If the use cases involve mapping attributes from authentication sources or PCV instances to the access tokens, runtime procedures obtain attribute values for each token request, but persistent grants do not store with attributes or their values.

Persistent grants, and any associated attributes and their values, remain valid until the grants expire or until PingFederate explicitly revokes them or cleans them up. PingFederate's persistent grant cleanup routine manages expired grants based on the Persistent Grant Max Lifetime policy setting.



Note:

PingFederate does not factor in the **Persistent Grant Idle Timeout** setting during grant cleanup. Ensure the grant datastore has the disk space needed to store expired grants because they exceeded the Persistent Grant Idle Timeout setting.

Managing expired persistent grants

PingFederate removes expired persistent grants once a day. The cleanup task removes 500 expired grants at a time until all expired grants are removed.

About this task

If expired grants are growing rapidly, you can optionally increase the frequency of the cleanup task.



Increasing the frequency of the cleanup task or the number of expired sessions to be removed per batch adds more workload to your storage server. Make gradual changes, if any, to observe the impact.



Important:

In a clustered PingFederate environment, the cleanup task runs only on the console node. If adjustments are required, make them on the console node. No changes are required on any of the engine nodes.

When storing persistent grants on a PingDirectory server that is version 7.0 or later, you can use the PingFederate cleanup task or configure a cleanup plugin in PingDirectory instead. The plugin allows finegrained control over various aspects of the cleanup task, which might improve the performance impact. For more information and configuration steps, see Managing expired persistent grants in PingDirectory on page 828.

Steps

- 1. Optional: Adjust the frequency of the cleanup task.
 - a. Edit the timer-intervals.xml<pf install>/pingfederate/server/default/data/ config-store directory.
 - b. Update the AccessGrantCleanerInterval value, in milliseconds.
 - The default value is 86400000, which is 24 hours.
 - c. Save your changes.
- 2. Optional: Adjust the number of expired grants to be removed per batch.
 - a. Edit the configuration file relevant to your storage platform.

This configuration file is located in the <pf install>/pingfederate/server/default/ data/config-store directory, as described in the following table.

Storage platform	Configuration file	
Database server	org.sourceid.oauth20.token.AccessGra	ntManager
PingDirectory	org.sourceid.oauth20.token.AccessGra	ntManager
Microsoft Active Directory	org.sourceid.oauth20.token.AccessGra	ntManager
Oracle Directory Server Enterprise Edition or Oracle Unified Directory	org.sourceid.oauth20.token.AccessGra	ntManager

b. Update the ExpiredGrantBatchSize value.

The following example shows an updated value of 400.

```
file, located in the<?xml version="1.0" encoding="UTF-8"?>
<c:config xmlns:c="http://www.sourceid.org/2004/05/config">
   <c:item name="ExpiredGrantBatchSize">400</c:item>
```

The default value is 500.

- c. Save your changes.
- 3. After you have made changes, restart PingFederate.

In a clustered PingFederate environment, you do not have to change or restart PingFederate on any of the engine nodes.

Managing expired persistent grants in PingDirectory

When storing OAuth persistent grants on a PingDirectory server that is version 7.0 or later, you can configure a cleanup plugin in PingDirectory to remove expired data from your directory server.

About this task

This PingDirectory plugin allows fine-grained control over various aspects of the cleanup task. For example, you can configure the maximum number of updates per second to improve the performance impact.

Steps

1. Disable the PingFederate cleanup task.



Important:

For a clustered PingFederate environment, make these change on the console node. No changes are required on any of the engine nodes.

- a. Edit the <pf install>/pingfederate/server/default/data/config-store/timerintervals.xml file.
- b. Update the AccessGrantCleanerInterval value to 0.
- c. Save your changes.
- d. Restart PingFederate.
- 2. Configure an instance of the PingDirectory plugin to clean up expired data.
 - a. Sign on to the PingDirectory administrative console.
 - b. Go to Configuration # Plugin Root.
 - c. Click New Plugin and then select Clean up Expired PingFederate Persistent Access Grants Plugin.
 - d. Configure a new instance of the Clean up Expired PingFederate Persistent Access Grants Plugin.

See the following table for information about each required field.

Field	Description
Name	The name of this plugin instance.
Enabled	The status of this plugin instance.
	Select the check box to enable this plugin instance. Clear the check box to disable this plugin instance.
	This check box is not selected by default.

e. Click Save.

Managing cleanup of persistent grants

PingFederate is capable of capping the number of persistent grants based on a combination of user, client, grant type, and authentication context.

About this task

Capping the number of persistent grants helps limit the data stored for persistent grants, especially in scenarios where clients frequently request authorization in a single context.

When PingFederate needs to record a new grant, it checks whether such creation will push the number of grants beyond the limit. If it does, PingFederate creates the grant and then removes just enough grants so that the number of grants is capped at the limit. This cleanup task starts from the oldest grant, expired or not, and continues forward if it needs to remove multiple grants. For performance reasons, this cleanup task also limits the number of grants it can remove per attempt. If it cannot remove all grants in excess of the limit, it removes what it can and repeats the process when PingFederate needs to record a new grant.

This cleanup runs on every engine node in a clustered PingFederate environment. Also, it does not replace the cleanup task or the PingDirectory plugin engineered to manage expired grants. Working together, they keep the size of the grant datastore under control.

The default limit is 100 grants per user, client, grant type, and authentication context. Depending on the storage platform, the default maximum number of grants that this cleanup task can remove per attempt varies.

This cleanup task is enabled on new installations. When upgrading from version 9.1 or an earlier version, it is disabled. You can enable it by editing an XML configuration file.

Steps

1. Edit the configuration file relevant to your storage platform.

This configuration file is located in the $< pf_install > / pingfederate / server / default / data / config-store directory, as described in the following table.$

Storage platform	Configuration file	
Database server	org.sourceid.oauth20.token.AccessGran	ManagerJd

Storage platform	Configuration file	
PingDirectory	org.sourceid.oauth20.token.AccessGran	- ManagerLD
Microsoft Active Directory	org.sourceid.oauth20.token.AccessGran	- ManagerLD
Oracle Directory Server Enterprise Edition or Oracle Unified Directory	org.sourceid.oauth20.token.AccessGran	tManagerLD

2. Locate for the following comments.

```
. . .
<!--
   Maximum number of persistent grants allowed to store in the database
per user,
    client and grant type and authentication context qualifier.
    Setting this to a value <= 0 will turn this limit off
    Default configuration:
    <c:item name="maxPersistentGrants">100</c:item>
<c:item name="maxPersistentGrants">100</c:item>
    Maximum number of persistent grants to delete when max allowed is
reached
    during new grant creation.
    Setting this to a value <= 0 will turn this limit off
    Default configuration:
    <c:item name="maxPersistentGrantsToRemoveBatchSize">n</c:item>
<c:item name="maxPersistentGrantsToRemoveBatchSize">n/c:item>
```

The maxPersistentGrants value represents the maximum number of grants based on a combination of user, client, grant type, and authentication context.

The maxPersistentGrantsToRemoveBatchSize value represents the maximum number of grants that the cleanup task would remove per attempt. Its default value (n) varies depending on the storage platform, 50 for a database server and 10 for a directory server.



Note:

The maxPersistentGrants and maxPersistentGrantsToRemoveBatchSize items exist only on new installations starting with version 9.2. When upgrading from version 9.1 or an earlier version, the upgrade tools only insert the comments for reference.

3. Optional: Adjust the maxPersistentGrants and maxPersistentGrantsToRemoveBatchSize values.

Use integers only.

4. To enable this cleanup task after upgrading from version 9.1 or an earlier version, insert the maxPersistentGrants and maxPersistentGrantsToRemoveBatchSize items into the configuration file.

You can use the default values based on the inline comment. You can also adjust the values to suit the needs of your organization.

5. Save your changes.

6. Restart PingFederate.

For a clustered PingFederate environment, perform these steps on the console node, and then click **Replicate Configuration** on **System # Server # Cluster Management**. You do not have to restart PingFederate on any running engine node.

Specifying the domain of the PF cookie

PingFederate identifies sessions by their respective PingFederate cookie. You can specify the domain of these cookies.

About this task

By default, the PingFederate cookie is set without domain information in the HTTP header.

```
Set-Cookie: PF=zOv4xxmzDI2rx1TFBFy78X;Path=/;Secure;HttpOnly
```

You can configure PingFederate to return the Set-Cookie HTTP header with domain information, as needed.

Steps

- 1. Edit the cookie-config.xml file.
- 2. Modify the cookie-domain element.

Example:

```
<c:item name="cookie-domain">.example.com</c:item>
```

- 3. Save the change.
- 4. Restart PingFederate.

For a clustered PingFederate environment, perform these steps on the console node, and then click **Replicate Configuration** on **System # Server # Cluster Management**. You do not have to restart PingFederate on any running engine node.

Result

After you activate this change, PingFederate includes domain information in its Set-Cookie HTTP header.

```
Set-Cookie:
```

PF=aDfPx6uwbbWGFhwE6zEhEG; Path=/; Domain=.example.com; Secure; HttpOnly

Specifying the domain of the PF.PERSISTENT cookie

PingFederate identifies persistent authentication sessions by their respective PF.PERSISTENT cookie. You can specify the domain of this cookie.

About this task

By default, the PF.PERSISTENT cookie is set without domain information in the HTTP header.

```
Set-Cookie: PF.PERSISTENT=UoBlPlf16V2oYAEPot2DnpUOXxitK7au;Path=/;Expires=Sat, 06-Nov-2021 00:48:08 GMT;Max-Age=94608000;Secure;HttpOnly
```

You can configure PingFederate to return the Set-Cookie HTTP header with domain information, as needed.

Steps

1. Edit the <pf_install>/pingfederate/server/default/data/config-store/persistent-session-cookie-config.xml file.

2. Modify the cookie-domain element.

Example:

<c:item name="cookie-domain">.example.com</c:item>

- 3. Save the change.
- 4. Restart PingFederate.

For a clustered PingFederate environment, perform these steps on the console node, and then click Replicate Configuration on System # Server # Cluster Management. You do not have to restart PingFederate on any running engine node.

Result

After you activate this change, PingFederate includes domain information in its Set-Cookie HTTP header.

```
Set-Cookie:
```

PF.PERSISTENT=tOYwPM7VFMeluUyeu0EKQLL0DCJyVOqG;Path=/;Domain=.example.com;Expires=Sat, 06-Nov-2021 01:00:34 GMT; Max-Age=94608000; Secure; HttpOnly

Extending the lifetime of the PingFederate cookie

PingFederate identifies sessions by their respective PingFederate cookies. You can manually extend the lifetime of these cookies.

About this task

Some adapters, such as the HTML Form Adapter, also utilize the PingFederate cookie to manage their adapter-sessions. The PingFederate cookie is a session cookie by default. You can extend the lifetime of the PingFederate cookie by making it a persistent cookie. Unlike session cookies, persistent cookies are saved to disk, enabling the browser to reuse them when restarted.



Alternatively, you can configure PingFederate to store authentication sessions externally and leverage them as users request protected resources after restarting their browsers. For more information, see Sessions on page 276.

Steps

- 1. Edit the session-cookie-config.xml file, located in the <pf install>/pingfederate/ server/default/data/config-store directory.
- Modify the cookie-max-age value.

The default value, -1, makes the PingFederate cookie a session cookie. A positive integer defines the age of the persistent cookie in seconds.

- 3. Save the change.
- 4. Restart PingFederate.

For a clustered PingFederate environment, perform these steps on the console node, and then click Replicate Configuration on System # Server # Cluster Management. You do not have to restart PingFederate on any running engine node.

Configuring forward proxy server settings

You can configure PingFederate to send web traffic, such as HTTP and HTTPS, so that it initiates through a forward proxy server.

Steps

- 1. Edit the <pf install>/pingfederate/bin/run.properties file.
- 2. Locate for the following properties.

```
#http.proxyHost=<HTTP_PROXY_HOST>
#http.proxyPort=<HTTP_PROXY_PORT>
#https.proxyHost=<HTTPS_PROXY_HOST>
#https.proxyPort=<HTTPS_PROXY_PORT>
#http.nonProxyHosts=*.internal.com|localhost
```

- 3. Optional: Configure forward proxy server settings for HTTP traffic.
 - a. Remove the number sign (#) in front of http.proxyHost and http.proxyPort.
 - b. Enter the hostname or the IP address of the forward proxy server.
- 4. Optional: Configure forward proxy server settings for HTTPS traffic.
 - a. Remove the number sign in front of https.proxyHost and https.proxyPort.
 - b. Enter the hostname or the IP address of the forward proxy server.
- 5. Optional: Configure an exclusion list.
 - a. Remove the number sign in front of http.nonProxyHosts.
 - b. Specify one or more destinations where PingFederate is not required to proxy its HTTP and HTTPS traffic through the forward proxy server.

This property supports multiple values separated by the pipe character (|) and the wildcard character (#) for pattern matching.

```
*.example.com|localhost
```

- 6. Save your changes.
- 7. Restart PingFederate.

For a clustered PingFederate environment, repeat these steps on each node.

Adding custom HTTP response headers

The PingFederate administrative console and runtime server are capable of returning custom HTTP response headers, such as HTTP Strict-Transport-Security (HSTS), to enforce HTTPS-based access and P3P for Microsoft Internet Explorer interoperability.

Steps

- 1. Edit the response-header-admin-config.xml file or the response-header-runtime-config.xml file, or both, located in the cpf_install>/pingfederate/server/default/data/config-store directory.
- 2. Save your changes.
- 3. Restart PingFederate.

For a clustered PingFederate environment, perform these steps on the console node, and then click **Replicate Configuration** on **System # Server # Cluster Management**. You do not have to restart PingFederate on any running engine node.

Configuring validation for the AudienceRestriction element

You can configure validation for the AudienceRestriction value in a SAML response.

About this task

For any identity provider (IdP) connection configured with multiple virtual server IDs, the **AudienceRestriction** value in a SAML response must match the virtual server ID information embedded in the protocol endpoint at which PingFederate receives the message.

You can disregard this validation condition on a per-connection basis.

Steps

- 2. Optionally, if you want to disregard the validation condition for an IdP connection, add its **Partner's Entity ID** value as an entry inside the c:map element.

Example:

Result:

In this example, the first entry adds the IdP connection with a **Partner's Entity ID** of www.example.com to the list. This is so that PingFederate no longer returns an error if the **AudienceRestriction** value in a SAML response does not match the virtual server ID information embedded in the protocol endpoint at which PingFederate receives the message. The second entry has the same effect for the IdP connection with a **Partner's Entity ID** of www.example.org.

- 3. Save your changes.
- 4. Restart PingFederate.

For a clustered PingFederate environment, perform these steps on the console node, and then click **Replicate Configuration** on **System # Server # Cluster Management**. You do not have to restart PingFederate on any running engine node.

Customizing the OpenID Provider configuration endpoint response

The OpenID Provider (OP) configuration endpoint at /.well-known/openid-configuration provides configuration information for the OAuth clients to interface with PingFederate using the OpenID Connect (OIDC) protocol.

About this task

As needed, you can customize the amount of configuration information by modifying a template file. You can also add conditional statements to return different responses, based on information from the requests to suit multiple use cases simultaneously.

Steps

1. Edit the openid-configuration.template.json file, located in the <pf install>/ pingfederate/server/default/conf/template directory, to specify the desired information to be returned by the OP configuration endpoint.

Multiple samples are provided, including sample statements using the \$HttpServletRequest and \$HttpServletResponse objects to get and set values.

2. Save your changes.

Template customization does not require a restart of PingFederate. For a clustered PingFederate environment, repeat these steps on each node.

Customizing the heartbeat message

The heartbeat endpoint, /pf/heartbeat.ping, returns a customizable OK browser message and an HTTP 200 status indication if the PingFederate server is running.

About this task

You can customize the message by modifying the pf.heartbeat.system.monitoring PingFederate property.



Note:

If a GET request receives a connection error or an HTTP status code other than 200, the server associated with the endpoint is down or malfunctioning.

Steps

1. Set the pf.heartbeat.system.monitoring property in the <pf install>/pingfederate/ bin/run.properties file to true or false.

When pf.heartbeat.system.monitoring is set to false, the /pf/heartbeat.ping endpoint returns OK. When set to true, the /pf/heartbeat.ping endpoint returns all available stats.

- 2. Restart PingFederate.
- 3. If you want to customize the information returned by the heartbeat endpoint, edit the heartbeat.page.template file, located in the <pf install>/pingfederate/server/ default/conf/template directory.

Template customization does not require a restart of PingFederate. For a clustered PingFederate environment, repeat these steps on each node.

- 4. If you want to specify percentiles in addition to, or in place of, the default 90th percentile in the statistics reported on the heartbeat:
 - a. Edit the com.pingidentity.monitoring.MonitoringService.xml file located in the <pf install>/pingfederate/server/default/data/config-store directory.
 - b. Change the value of the StatisticsPercentilesList item to the preferred percentile. You can enter a single value, such as 99.9, or multiple values separated by commas, such as 80,90,99.5.



Note:

The StatisticsPercentilesList item allows you to customize the percentiles displayed in the heartbeat endpoint response of the metrics ending in .<StatisticsPercentilesList>.percentile. Percentiles can be a helpful way to understand a metric's distribution and identify patterns or trends over time. They can also be used to set performance targets or to identify bottlenecks in a system.

In the context of server response metrics, you can use percentiles to compare a server's response time to other servers or previous periods. For more information, including a complete list of server metrics and their descriptions, see *Liveliness and responsiveness*.

Example: Setting the StatisticsPercentilesList item to 50 will display the 50th percentile of a server's response time in the heartbeat endpoint response. A value of 200 milliseconds for the 50th percentile means that 50% of the server's responses were faster than 200 milliseconds, while 50% were slower. Similarly, a value of 500 milliseconds for the 95th percentile means that 95% of the server's responses were faster than 500 milliseconds, while 5% were slower.

c. Save your changes.

Customizing the favicon for application and protocol endpoints

PingFederate provides a favorite icon (favicon) for its application, such as /idp/startSSO.ping, and protocol endpoints, such as /idp/SSO.saml2.

Steps

- 1. In the <pf install>/pingfederate/server/default/conf/template/assets/images directory, replace the favicon.ico file.
- 2. Restart PingFederate. For a clustered PingFederate environment, repeat these steps on each node.

Configuring the behavior of searching multiple datastores with one mapping

If a datastore uses results from previous queries as input, and if the previous queries return no result, PingFederate records a warning message in the server log and continues with the request by querying the next datastore in the attribute source setup.

About this task

This default behavior applies to all lookup configurations using multiple datastores in one mapping. For more information, see Attribute mapping with multiple data sources on page 362.

If you prefer PingFederate to abort the request immediately, which is the default behavior of many earlier versions of PingFederate, you can override the behavior by modifying a configuration file. Like the default behavior, this override also applies to all lookup configurations using multiple datastores in one mapping.

Steps

1. Edit the org.sourceid.saml20.domain.AttributeMapping.xml file, located in the <pf install>/pingfederate/server/default/data/config-store directory.



If this file does not exist, you must create it.

2. To override the default behavior, change the value of the AbortOnAttrLookupFailure element from false, the default value, to true.

The following is an example of a modified

org.sourceid.saml20.domain.AttributeMapping.xml file.

```
<?xml version="1.0" encoding="UTF-8"?>
<c:config xmlns:c="http://www.sourceid.org/2004/05/config">
    <c:item name="AbortOnAttrLookupFailure">true</c:item>
```



Note:

Removing the org.sourceid.saml20.domain.AttributeMapping.xml file from the <pf install>/pingfederate/server/default/data/config-store directory also has the same effect as setting the value of the AbortOnAttrLookupFailure element to true.

For a clustered PingFederate environment, perform these steps on the console node, and then click Replicate Configuration on System # Server # Cluster Management.

Example

Expected result when this override is set

If a datastore uses results from previous queries as input, and if the previous queries return no result, PingFederate records an error message in the server log, aborts the request immediately, and returns an error message to the user, the application, or the partner.

System settings

The **System** menu provides access to system-related settings.

Depending on the setup of PingFederate, menu items vary. Menu items are organized into four groups:

- Server
- Metadata
- Monitoring & Notifications
- External Systems

Server

You can configure various PingFederate settings.

On the Extended Properties window, you can configure protocol settings and local administrative accounts, import your license file, export or import a configuration archive, review other runtime servers in the same cluster and replicate configuration from the administrative server to the clustered runtime servers. You can also configure virtual host names, and configure extended properties for connections and OAuth clients. See the following topics for more information:

- Administrative accounts on page 842
- License management on page 847
- Configuration archive on page 851
- Cluster management on page 854
- Virtual host names on page 854
- Extended properties on page 856

Protocol settings

On the Protocol Settings window, you can configure the base URL of your PingFederate environment, and settings for WS-Trust STS authentication, outbound provisioning, and IdP discovery.

Federation information identifies your federation deployment to your partners, according to the protocols you support.

About this task

You must provide an ID that uniquely identifies your federation gateway for each protocol you support. For WS-Trust security token service (STS), IDs are required for both SAML 2.0 and SAML 1.x, regardless of browser-based single sign-on (SSO) protocol support or the type of token expected to be issued, to ensure that the STS will perform correctly under all conditions.



Each ID normally applies across all connection partners for a given protocol. However, if your implementation requires different IDs for the same protocol, you can use virtual server IDs. For more information, see Federation planning checklist on page 104.

Steps

- 1. Go to System # Server to open the Protocol Settings window.
- 2. On the **Federation Info** tab, provide the required information.

For more information, see the following table.

Field	Description
Base URL	The fully qualified host name, port, and path (if applicable) on which the PingFederate server runs. This field is used to populate configuration settings in metadata files. For more information, see <i>Metadata export</i> on page 860.
SAML 2.0 Entity ID	This ID defines your organization as the entity operating the server for SAML 2.0 transactions. It is usually defined as an organization's URL or a DNS address, for example: pingidentity.com. The SAML SourceID used for artifact resolution is derived from this ID using SHA1.
SAML 1.x Issuer/ Audience	This ID identifies your federation server for SAML 1.x transactions. As with SAML 2.0, it is usually defined as an organization's URL or a DNS address. The SourceID used for artifact resolution is derived from this ID using SHA1.
SAML 1.x Source ID	(Optional) If supplied, the Source ID value entered here is used for SAML 1.x, instead of being derived from the SAML 1.x Issuer/Audience.
WS-Federation Realm	The URI of the realm associated with the PingFederate server. A realm represents a single unit of security administration or trust.

3. Click **Next** and continue with the rest of the configuration.



When editing an existing configuration, you can also click Save as soon as the administrative console offers the opportunity to do so.

Configuring WS-Trust settings

You can configure PingFederate to require that client applications provide credentials to access the security token service (STS).

About this task

While this is an optional configuration, it is recommended for identity provider (IdP) configurations using the Username Token Processor. For other token processors and token generators, trust in the identity of the client is conveyed within the token itself and verified as part of processing. However, you can still configure authentication requirements to add another layer of security by limiting access to only authenticated clients.



You can configure STS authentication to either apply globally to all token formats and for all IdP and service provider (SP) partner connections, or token-to-token mappings, using more fine-grained controls at the connection level through issuance criteria.

Steps

- 1. Go to System # Serverto open the Protocol Settings window
- 2. On the WS-Trust STS Settings tab, click Configure WS-Trust STS Authentication.

Follow the configuration wizard to complete the task. For more information, see Configuring STS authentication on page 951

3. Click **Next** and continue with the rest of the configuration.



Tip:

When editing an existing configuration, you can also click **Save** as soon as the administrative console offers the opportunity to do so.

Configuring outbound provisioning settings

Select the database that PingFederate should use internally to facilitate provisioning for service providers when PingFederate is configured as an identity provider (IdP).

Before you begin

Before configuring outbound provisioning settings, you must enable outbound provisioning through the pf.provisioner.mode property in the <pf install>/pingfederate/bin/run.properties file. For more information, see *Configuring PingFederate properties* on page 752.

If you want to use failover provisioning, configure the provisioner.node.id and provisioner.failover.grace.period properties, which are also located in <pf install>/ pingfederate/bin/run.properties. These properties are described in Deploying provisioning failover on page 211.

About this task

The database stores the state of synchronization between the source datastore and the target datastore, enabling periodic checking to determine whether updates are required at the target site. PingFederate checks the source datastore for changes every minute by default. As needed, you can change the provisioning synchronization frequency on this tab as well.



CAUTION:

Use the built-in HSQLDB only for trial or training environments. For testing and production environments, always use a secured external storage solution for proper functioning in a clustered environment.

Testing involving HSQLDB is not a valid test. In both testing and production, it might cause various problems due to its limitations and HSQLDB involved cases are not supported by Pingldentity.



Note:

PingFederate is tested with Amazon Aurora (MySQL and PostgreSQL), Microsoft SQL Server, Oracle Database, Oracle MySQL, and PostgreSQL as internal provisioning datastores. A demonstrationonly, embedded HSQLDB database is installed by default. Scripts to aid setup are in the directory <pf install>/pingfederate/server/default/conf/provisioner/sgl-scripts.



CAUTION:

A pre-installed, default HSQLDB database is selected for initial setup and testing. However, we strongly recommend that you use the built-in HSQLDB only for trial or training environments. For testing and production environments, always use a secured external storage solution for proper functioning in a clustered environment.

Testing involving HSQLDB is not a valid test. In both testing and production, it may cause various problems due to its limitations and HSQLDB involved cases are not supported by Pingldentity.

Steps

- 1. Go to System # Server # Protocol Settings.
- 2. On the Outbound Provisioning tab, from the Internal Provisioning Data Store list, select a datastore.

If the datastore you want is not shown in the list, PingFederate is not yet configured to access the store. Click Manage Data Stores to create a connection to the datastore.

3. Change the **Synchronization Frequency** value.

The default value is 60 seconds.

4. Click **Next** and continue with the rest of the configuration.



When editing an existing configuration, you can also click **Save** as soon as the administrative console offers the opportunity to do so.

Configuring standard IdP Discovery

SAML 2.0 identity provider (IdP) Discovery provides a cookie-based look-up mechanism to identify a user's IdP dynamically during a service provider (SP)-initiated single sign-on (SSO) event when the IdP is not otherwise specified.

About this task

This mechanism can be helpful in cases where an SP might be a hub for several IdPs in an identity federation.

In addition to supporting SAML 2.0 IdP Discovery, PingFederate provides a cross-protocol, proprietary mechanism allowing a PingFederate SP server to write a persistent browser cookie. The cookie contains a reference to the IdP partner with whom the user previously authenticated for SSO. For more information, see Configuring IdP discovery using a persistent cookie on page 751.

Tip:

An SP can also include the discovery mechanism within the application. For instance, an SP can provide vanity URLs to isolate one set of end users from the others based on the URL of the requested resources. Another possible solution is to provide a user interface for the end users to enter information about their identity providers. With this approach, the application can start an SP-initiated SSO request with information about the IdP.

In the standard scenario, when a user requests access to a protected resource on the SP, commondomain browser cookies are used to determine where a user has authenticated in the past. Using this information, a PingFederate server can determine which IdP connection to use for sending an authentication request.

As an IdP Discovery provider, PingFederate can serve in up to three different roles: common domain server, common domain cookie writer, and common domain cookie reader. Each of these roles is necessary to support IdP Discovery. The roles can be distributed across multiple servers at different sites.

Common domain server

In this role the PingFederate server hosts a domain that its federation partners share in common. The common domain server allows partners to manipulate browser cookies that exist within that common domain. PingFederate can serve in this role exclusively or as part of either an IdP or an SP federation role, or both.

Common domain cookie writer

When PingFederate is acting in an IdP role and authenticates a user, it can write an entry in the common domain cookie, including its federation entity ID. An SP can look up this information on the common domain, not the same location as the common domain server described above.

Common domain cookie reader

When PingFederate is acting as an SP and needs to determine the IdPs with whom the user has authenticated in the past, it reads the common domain cookie. Based on the information contained in the cookie, PingFederate can then initiate an SSO authentication request using the correct IdP connection.

Steps

- 1. Go to the **Protocol Settings** window.
- 2. On the IdP Discovery tab, click Configure IdP Discovery.
- 3. On the **Domain Cookie Settings** tab, choose the discovery role or roles of your PingFederate server.
- 4. On the **Common Domain Service** tab, configure as follows.

Field	Description
Base URL of the	Enter the base URL of the PingFederate common domain service.
PingFederate Common Domain Service	A common domain service is where PingFederate reads or writes authentication information contained in shared cookies, as determined by whether your site is an SP or IdP, respectively. The service is shared if your PingFederate server is acting in both roles. You must use HTTPS for the common domain.
Pass Phrase and Confirm Pass	Enter and confirm the pass phrase that web applications must use to access the domain.

5. On the Local Common Domain Server tab, configure the required settings.

A local common domain server is where PingFederate reads (as an SP) or writes (as an IdP) a common domain cookie (CDC) for IdP Discovery.

Field	Description	
Common Domain	Enter the common domain.	
	Your entry must include an initial period (.), as in the following example.	
	.example.com	
Cookie Lifetime (Days)	Enter the lifetime of the CDC in days. The range is 1 to 1825 days. To indicate a non-persistent session cookie, enter -1 .	
Pass Phrase and Confirm Pass	Enter and confirm the pass phrase that web applications must use to access the domain.	

6. On the **Summary** tab, review and modify settings as needed. Then click **Save**.

Result:

The administrative console brings you back to the **IdP Discovery** tab.

7. Click **Next** and continue with the rest of the configuration.



When editing an existing configuration, you can also click Save as soon as the administrative console offers the opportunity to do so.

- 8. Perform one of the following actions to enable the setting of the common domain cookie at runtime: Choose from:
 - Make sure that, prior to launching any SSO events, the web application that implements IdP Discovery sets the cookie using the /idp/writecdc.ping application endpoint intended for that
 - Enable setting the cookie at runtime during SSO events by selecting the IdP Discovery check box on the Connection Options tab for the desired SP connection.

Reviewing protocol settings

You can review your protocol settings before saving them.

Steps

- To amend your configuration, click the corresponding tab title and then follow the steps to complete the task.
- To keep your changes, click Save.
- To discard your changes, click Cancel.

Administrative accounts

The PingFederate administrative console supports five authentication schemes.

The authentication schemes are:

- Native authentication
- LDAP authentication
- RADIUS authentication
- Certificate-based authentication
- OIDC-based authentication

PingFederate User Access Control

Account type	Administrative role	Access privileges
Admin	User Admin	Create users, deactivate users, change or reset passwords, and install replacement license keys.
Admin	Admin	Configure partner connections and most system settings, except the management of local accounts and the handling of local keys and certificates.
Admin	Expression Admin	Map user attributes by using the expression language, Object-Graph Navigation Language (OGNL).
		Important:
		Only Administrative users who have both the Admin role and the Expression Admin role:
		 Can be granted the User Admin role. This restriction prevents non-Expression Admin users from granting themselves the Expression Admin Role.
		• Can be granted write access to the file system or directory where PingFederate is installed. This restriction prevents a non-Expression Admin user from placing a data.zip file containing expressions into the <pre><pre>cypr_install</pre>/pingfederate/server/ default/deploy directory, which would introduce expressions into PingFederate.</pre>
Admin	Crypto Admin	Manage local keys and certificates.
Auditor	Not applicable	View-only permissions for all administrative functions. When the Auditor role is assigned, no other administrative roles can be set.



All four administrative roles are required to access and make changes through the following services:

- The /bulk, /configArchive, and /configStore administrative API endpoints
- The Configuration Archive window, accessed from System # Server, in the administrative console
- The Connection Management configuration item on the Service Authentication window, accessed from Security # System Integration

For native authentication, access and authorization are controlled by the local accounts defined on the Administrative Accounts window.

As needed, you can switch from native authentication to an alternative console authentication. Access and authorization are defined in the respective configuration file.

An administrative user can sign on from more than one browser or location. Moreover, multiple administrative users can sign on to the PingFederate administrative console at a time. You can optionally restrict the administrative console to one administrative user at a time by modifying the pf.console.login.mode property in the <pf install>/pingfederate/bin/run.properties file. Regardless of the property configuration, any number of auditors can sign on at any time.

Note:

For security, after three failed sign-on attempts from the same location within a short time period, the administrative console and the administrative API will temporarily lock out further attempts by the same user. The user must wait one minute to try again.

Local accounts defined on the Administrative Accounts window are shared between the administrative console and the administrative API if they are both configured to use native authentication, the default. If the administrative console is configured to use an alternative console authentication, the Administrative Accounts window appears only if the administrative API is left to use native authentication, and vice versa.



If you have connected PingFederate to PingOne for Enterprise, you can also single sign-on from the PingOne admin portal to the administrative console.

Enabling native authentication for the administrative console

When the administrative console is protected by native authentication, access is restricted to the local accounts you have defined.

Steps

- 1. Go to System # Serverto open the Administrative Accounts window.
- 2. In the <pf insall>/pingfederate/bin/run.properties file, change the value of the pf.console.authentication property to pf.console.authentication=native
- 3. Start or restart PingFederate.

Managing local accounts and role assignments

You can create, modify, update, or deactivate accounts in the Administrative Accounts window.

Steps

1. Go to System # Server # Administrative Accounts, and then perform any of the following actions.

Task	Steps
Create a local account	On the Administrative Accounts window, click Create User.
	 On the User Information tab, enter a username and other optional information.
	Note:
	If you want PingFederate to notify the user about password changes via email, you must supply an email address.
	c. On the Password Generation tab, enter a password or click Generate one-time

Task	Steps
	password to generate a random password for the account.
	Note:
	Upon successful authentication, the user will be required to change the password of the account immediately.
	d. On the Summary tab, review your configuration, modify as needed, and then click Done .
	 e. On the Administrative Accounts window, select the applicable account type, Auditor or Admin, and one or more administrative roles for an Admin account. f. Repeat these steps to create additional
	accounts.
Modify user information	On the Administrative Accounts window, select the account by its username.
	Note: Applicable only to active accounts.
	b. On the User Information window, update the record, and then click Done.c. Repeat these steps to update other accounts.
Update role assignments	 a. Select a different account type, Auditor or Admin, for one or more accounts. b. Select or clear the check boxes that correspond to the three administrative roles, User Admin, Admin, and Crypto Admin for one or more accounts.
	Note: Applicable only to the Admin accounts.
Deactivate or reactive a native	a. Click Deactivate or Activate under Action.b. Repeat this step to deactivate or reactive other accounts.
	Note: For traceability and accountability purposes, local accounts cannot be deleted Their

Task	Steps
	records are retained and they can be reactivated if needed.

2. To keep your configuration, click Save.

Enabling notification messages for account management events

Administrators can optionally enable notifications for account management events. If enabled, PingFederate generates notification messages based on the following events.

About this task

See the following table for a description of an event and the corresponding alert.

Event	Alert
An administrator	PingFederate generates a notification message to all administrators.
turns off the Notify Administrator of Account Changes option.	The message includes the username of the administrator who made the change.
An administrator's email address is updated by	PingFederate generates a notification message to the previous email address and another notification to the new email address.
another administrator.	The message includes the username of the administrator who made the change.
An administrator's password is changed.	PingFederate generates a notification to the administrator whose password has been changed.
	The message includes the username of the administrator who made the change.



Note:

Account management events are only applicable when native authentication is enabled for the administrative console, the administrative API, or both in the <pf install>/pingfederate/bin/ run.properties file. If you are using an alternative console authentication, notifications, if any, such as password changes, are handled by the third-party system.

Steps

- 1. Go to System # Serverto open the Administrative Accounts window.
- 2. Select the **Notify Administrator of Account Change** check box.

An email address must be provided for the applicable accounts.

- 3. Select a notification publisher instance from the list.
 - If you have not yet configured the desired notification publisher instance, click Manage Notification Publishers.
- 4. To keep your configuration, click **Save**.

User administrators can generate temporary passwords as they create new local accounts for new users on the Password Generation tab.

About this task

User administrators can also assign temporary passwords for existing users who forget their passwords. Upon successful authentication, the users are required to change their passwords immediately.

Steps

- 1. Go to System # Server # Administrative Accounts.
- 2. To generate a random password for the account, on the Password Generation tab, enter a password or click Generate one-time password, and then click Next.

Changing passwords

Administrative users and auditors can change the passwords of their local accounts at any time.

About this task



If you sign on to PingFederate using your network ID and password, you can change your password only at the network level. The new password will apply to PingFederate automatically the next time you log on.

Steps

- 1. Go to System # Server to open the Administrative Accounts window.
- 2. Click Change Password under Action for your account.
- 3. Enter your current password and new password twice in the related fields.



Important:

If you are the sole user administrator, take steps to ensure that you do not forget your new password.

4. To keep your configuration, click Save.

License management

PingFederate licensing is handled differently depending on whether you are installing and setting up PingFederate for the first time, or upgrading your existing PingFederate installation to a later version.

Initial PingFederate installations

During the initial setup for a new PingFederate installation:

- If you choose to connect to PingOne for Enterprise, PingFederate obtains and installs an evaluation PingFederate Bridge license from PingOne.
- If you choose not to connect to PingOne for Enterprise, you are prompted to upload a license file.

Depending on your licensing agreement, your PingFederate license might have an expiration date. If your license key is going to expire, or has expired recently, you can import a new license file to replace the existing license key through the administrative console.

The administrative console displays a warning message ahead of the expiry of your license. Optionally, you can configure PingFederate to notify the administrators ahead of the license expiration date.

If you choose to upgrade PingFederate using the Upgrade Utility, your current PingFederate license is automatically copied to the target installation if it is valid. If your current license is not valid, you must obtain a new license and specify its full path and filename when performing the upgrade.

If you choose not to use the Upgrade Utility, you must specify the full path and filename of a valid license when performing the upgrade.

Reviewing license information

You can find and review a summary of your PingFederate license on the administrative console.

Steps

- 1. Open the Help menu.
- 2. Click About.

Result:

The **License Summary** appears.



Important:

If the license specifies an expiration date, the license expires at the beginning of that day.

Requesting a new license key

You can request a new license key if needed.

Steps

- 1. Go to the Ping Identity licensing website.
- 2. Sign on, provide the required information, and then submit your request.

Result

You will hear from our licensing team as the team processes your request.

Installing a license key on a new or upgraded PingFederate server

You can install a license key on the PingFederate server from the administrative console.

Steps

- 1. Start PingFederate and sign on to the administrative console.
- 2. Go to System # Server to open the License Managementwindow, and then click Choose file to select the license file. Click Import.

Result:

If the license is for the wrong version of PingFederate or is found to be invalid for some other reason, PingFederate displays an error message.

When the license file is verified for use with your PingFederate instance, the license information is saved in the <pf install>/pingfederate/server/default/conf/pingfederate.lic file. 3. If you have a clustered PingFederate environment, go to System # Server to open the Cluster Management window, and then click Replicate Configuration.



Note:

You must use the Replicate Configuration window to initiate the transmission of the license file from the console node to all engine nodes. As an added measure, the administrative console reminds you to do so as well.

When an engine node receives the license information from the console node, it saves the new license information to the <pf install>/pingfederate/server/default/conf/pingfederate.lic license file.

For any engine node that was offline at the time of the import, when it restarts and joins the cluster, it consumes the new license information from the console node and applies the same processing logic.



PingFederate no longer maintains its license information in the install>/pingfederate/ server/default/data/.pingfederate.lic file, which is referred to as the secondary license file in the previous versions of PingFederate. The .pingfederate.lic file, if any, is ignored.

Installing a replacement license key

You can replace your existing PingFederate license key by importing a new license file using the PingFederate administrated console.

Steps

- 1. Start PingFederate and sign on to the administrative console.
- 2. Go to System # Serverto open the License Management window, and then click Choose File to select the license file. Click Import.

Result:

If the license is for the wrong version of PingFederate or is found to be invalid for some other reason, PingFederate displays an error message and keeps the existing license regardless of whether the existing license has expired.

If the new license does not include support for features found in your existing license, or if there is some other potential problem with the license, PingFederate advises and prompts you on whether to continue.

When the license file is verified for use with your PingFederate instance, the license information is saved in the <pf install>/pingfederate/server/default/conf/pingfederate.lic file.

The previous license file is renamed with a timestamp in the same conf directory.

3. If you have a clustered PingFederate environment, go to System # Server to open the Cluster Managementwindow, and then click Replicate Configuration.



Note:

You must use the Replicate Configuration window to initiate the transmission of the license file from the console node to all engine nodes.

When an engine node receives license information from the console node, if the issue date of the new license is not as recent as that of the existing license, the engine node ignores the new license from the console and logs the following warning message in the server log.

License from Console node ignored as Engine node has recently obtained license.



If you prefer to use the license from the console node, even the existing license on the engine node is more recent in terms of the issue date, manually remove or rename the <pf install>/ pingfederate/server/default/conf/pingfederate.lic license file on the engine node, then click Replicate Configuration on the Cluster Management window again.

If the issue date of the new license is more recent than or equal to that of the existing license, the engine node saves the new license information to the <pf install>/pingfederate/server/ default/conf/pingfederate.lic license file and activates it immediately. No restart is required.

For any engine node that was offline at the time of the import, when it restarts and joins the cluster, it consumes the new license information from the console node and applies the same processing logic.



Note:

PingFederate no longer maintains its license information in the install>/pingfederate/ server/default/data/.pingfederate.lic file, which is referred to as the secondary license file in the previous versions of PingFederate. The .pingfederate.lic file, if any, is ignored.

Configuring notification for licensing events

If your PingFederate license has an expiration date, you can configure PingFederate to notify the administrators ahead of the license expiration date to minimize potential service disruptions.

About this task

Steps

- 1. Go to System # Monitoring & Notifications to open the Runtime Notifications window.
- 2. Select the Notification for Servers Licensing Events check box.



Note:

This check box appears only if your PingFederate license has an expiration date.

- 3. Enter the email address of the intended recipient.
- 4. Select a notification publisher instance from the list.

If you have not yet configured the desired notification publisher instance, click Manage Notification Publishers.

Configuration archive

You can use configuration archives as backup files for the current PingFederate installation.

PingFederate automatically creates a time-stamped configuration (.zip) archive every time an administrator signs on to the administrative console and before an existing archive is imported. The archives are stored in the <pf install>/pingfederate/server/default/data/archive directory.

The automatic backup process typically completes without delays. For deployments with hundreds of connections or OAuth clients, or both, administrators can configure PingFederate to create configuration archives periodically instead.

Additionally, administrators can export the current configuration to a .zip file in the Configuration Archive window. This window is only available to administrators whose accounts have been assigned the User Admin, Admin, Crypto Admin, and Expression Admin roles.



The Expression Admin role must be assigned to give administrators sufficient permissions to create configuration archives.



Warning:

The backup file contains your complete configuration. To protect your data, confirm the backup file is protected with appropriate security controls in place before exporting it.

Sharing the archive is a security risk because the private keys are stored in the archive. An archive should only be shared if the security of that instance is not important, such as a development or test environment.

On the Configuration Archive window, administrators can import an existing archive for immediate deployment into a running PingFederate server.

Administrators can also deploy a configuration archive manually by copying the .zip file to the <pf install>/pingfederate/server/default/data/drop-in-deployer directory. After copying the .zip file, it must be renamed to data.zip.

Configuration archives are intended for administrative-console configuration only. The following files are not included in the archives:

- Launch scripts in the <pf install>/pingfederate/bin and <pf install>/pingfederate/ sbin directories.
- Web container configuration files in the <pf install>/pingfederate/etc directory.
- Log files in the <pf install>/pingfederate/log directory.
- Database drivers and program files from adapters and any other plugins in the <pf install>/ pingfederate/server/default/lib and <pf install>/pingfederate/server/default/ deploy directories.
- Other files, including the license file, the advanced cluster configuration files, and the user-facing email and HTML templates, in the <pf install>/pingfederate/server/default/conf directory.

If any changes have been made to files that are not part of the configuration archive, those files must be preserved manually.



Tip:



Important:

Draft connections in archives are not imported. Complete any unfinished partner connections if you want to include them in a full backup archive or in an archive to be used for configuration migration.

Configuring a backup schedule

For deployments with hundreds of connections or OAuth clients, or both, administrators can configure PingFederate to create configuration archives periodically.

Steps

Edit the <pf install>/pingfederate/server/default/data/config-store/ org.sourceid.saml20.domain.mgmt.impl.DataArchiveBackup.xml file. See the following table

Property	Description
ScheduledBackupEnabled	When set to true, PingFederate creates a configuration archive daily if configuration has changed since the creation of the last archive. This is true regardless of the backup method used.
	The default value is false.
BackupTime	The local time at which PingFederate creates a configuration archive when the ScheduledBackupEnabled property is set to true.
	Ignored when the ScheduledBackupEnabled is set to false.
	The default value is 00:00:00, which represents midnight.
NumTriesWithIntegrityFailure	Upon the creation of the scheduled configuration archive, if PingFederate detects a configuration change during the scheduled backup process, it retries immediately. The NumTriesWithIntegrityFailure property indicates the maximum number of attempts.
	Ignored when the ScheduledBackupEnabled is set to false.
	The default value is 3.
BackupOnAdminLogin	When set to false, PingFederate does not create a configuration archive when an administrator signs on. Regardless of the value of this property, PingFederate always creates a configuration archive before an archive is imported. The default value is true.

Property	Description
MaxOldArchiveFiles	The number of older configuration archives that PingFederate keeps. When the limit is reached, the oldest file is removed.
	The default value is 25 files.

The ScheduledBackupEnabled and BackupOnAdminLogin properties are not mutually exclusive. For example, if your deployment has 50 connections, you can enable both automatic and scheduled backup processes.



Note:

If you have a clustered PingFederate environment, edit the

org.sourceid.saml20.domain.mgmt.impl.DataArchiveBackup.xml file on the console node.

Exporting an archive

Administrators can export the current administrative-console configuration to a .zip file.

About this task

This process can be performed in the Configuration Archive window. This window is only available to administrators whose accounts have been assigned the User Admin, Admin, and Crypto Admin roles.

Steps

On the Export tab, click Export, then save the .zip file.



Warning:

The backup file contains your complete configuration. To protect your data, confirm the backup file is protected with appropriate security controls in place before exporting it.

Sharing the archive is a security risk because the private keys are stored in the archive. An archive should only be shared if the security of that instance is not important, such as a development or test environment.

Importing an archive

Administrators can import an administrative-console configuration from a .zip file.

About this task

This process is performed in the Configuration Archive window. This window is only available to administrators whose accounts have been assigned the User Admin, Admin, and Crypto Admin roles.

When an administrator initiates deployment of a configuration archive using the **Import** tab, PingFederate displays error messages if there are any missing plugin components, such as adapters, database drivers, or token translators, on which the archive depends, or any mismatches of PingFederate licensing authorization. The administrator can choose to force the deployment and then install the necessary files later.



Installation of any missing database drivers or other third-party libraries will require a restart of PingFederate.

CAUTION:

Deploying a configuration archive, either manually or by using the administrative console, always overwrites all existing configuration data.

Steps

- 1. On the **Import** tab, choose the desired configuration archive from your system.
- 2. Select the Force Import check box if you want PingFederate to deploy the archive regardless of whether dependency errors are detected.



Important:

If you make this selection, consult the server start-up console or the server log for any messages concerning missing plugin components or other errors.

3. Click Import.

Result:

The administrative console prompts you to confirm the import process.

Cluster management

When multiple PingFederate servers are set up to run as a cluster, the administrative console provides a Cluster Management window.

Whenever applicable changes are made through the administrative console, a message appears at the top of the console to serve as a reminder to go to the Cluster Management window and to replicate the current console configuration to all server nodes in the cluster.

As of PingFederate 10.0, when your cluster is running multiple maintenance versions, a second message is displayed saying that more than one version is being run. The Cluster Management window also contains a new Version column that displays the version of PingFederate being run on each server node.

Replicating configuration

You can replicate a configuration from the PingFederate administrative console.

Steps

- 1. Go to System # Server.
- 2. In the Cluster Management window, click Replicate Configuration.
- 3. Click Done.

Result

The replication process takes only a moment. Requests coming into the runtime servers continue to be processed while the updated configuration is being loaded into memory. When the reload is complete, the updated configuration is activated for new requests.

Virtual host names

You can optionally define a list of alternate domain names at which PingFederate receives application and protocol messages.

This is done in the Virtual Host Names window. When configured, PingFederate honors the originally requested host throughout all browser redirects and metadata retrieval if the requested host matches one of the virtual host names. This capability allows you to fully support any number of branded URLs regardless of configured use cases within a single PingFederate environment.

SAML inbound message

In certain contexts, the SAML specifications require that XML messages include a URL identifying the host name to which the sender directed the message. As the recipient of such messages, PingFederate validates that the value matches the location where the message is received, which is the Base URL value defined in the Protocol Settings window on the Federation Info tab.

When virtual host names are configured, PingFederate takes them into consideration as part of its message-security validation process, in addition to its base URL.

OAuth client authentication using the private key jwt client authentication method

An OAuth client can authenticate with an authorization server by presenting a signed JWT. Per specification, the client must include the intended recipient as the aud claim value in its JWT. When acting as the authorization server, PingFederate verifies that the destination of the aud claim value matches either its base URL or the Token Endpoint Base URL value defined in the Authorization Server Settings window.

When virtual host names are configured, PingFederate uses them in its verification process as well.



Virtual host names and virtual server IDs serve different purposes. The latter provides separate unique identifiers on a per-connection basis for a federation deployment, normally in the same domain. For more information, see Multiple virtual server IDs on page 106. Virtual host names and virtual server IDs are not mutually exclusive. Depending on your use cases and infrastructure, you can configure both virtual server IDs and virtual host names in your PingFederate environment.

Multiple site certificates

When multiple domain names are involved, you can configure PingFederate with multiple site certificates so that PingFederate can serve a different site certificate based on the requested host. For more information, see *Manage SSL server certificates* on page 601.

Configuring virtual host names

Use the PingFederate administrative console to configure and manage virtual host names.

Steps

- 1. Go to System # Server to open the Virtual Host Names window.
- 2. Perform any of the following actions.

Option	Action
Add a new entry	Enter the desired value and click Add
Modify an existing entry	Click Edit or Update
Remove an existing entry	Click Delete
Keep your changes	Click Save
Discard your changes	Click Cancel

Extended properties

You can add any number of extended properties to store additional information about connections, OAuth clients, or both.

You add these extended properties in the Extended Properties window as described in Defining extended properties on page 856. When adding an extended property, you can define it as a single-value property or a multivalued property. These extended properties become available to all connections and clients. As you create or update a connection or a client, you can populate values for any of them. For OAuth clients, if dynamic client registration is configured and enabled, developers can populate extended property values by including them in the client registrations.

Authentication policies

You can leverage extended properties to drive authentication experience and requirements by configuring an instance of the Extended Property Authentication Selector for each property that matters, placing this selector instance in an authentication policy, and defining a policy path for each selector result value. At runtime, PingFederate routes browser-based single sign-on (SSO) requests, OAuth authorization requests, and OAuth grant management requests to the desired authentication sources based on the applicable policy.

For more information, see Configuring the Extended Property Authentication Selector on page 225.

OAuth attributes fulfillment and issuance criteria

You can use extended properties as attribute sources when fulfilling persistent grants and token contracts. You can also define issuance criteria to verify extended property values.

Defining extended properties

You can define extended properties using the PingFederate administrative console.

Steps

- 1. Go to System # Server # Extended Properties.
- 2. Add the extended properties.
 - a. Enter the applicable property name under **Name**.
 - After being added and saved, the name cannot be modified, but you can delete the extended property completely.
 - b. Optional: In the **Description** field, enter a description for the extended property.
 - c. If the extended property should allow multiple values, select the Multivalued check box .



Note:

If you have initially added a single-valued extended property, you can make it a multivalued extended property by selecting the Multivalued check box later.

If you have initially added a multivalued extended property and you clear the **Multivalued** check box later, when you try to make changes to a connection or an OAuth client that has been configured with multiple values for this extended property, the administrative console prompts you to reconfigure the connection or the client until only one value exists for this extended property.

- d. Click Add.
- Optional: Repeat these steps to define additional extended properties.

Click Edit, Update, or Cancel to make or undo a change to an existing entry. Click Delete or **Undelete** to remove an existing entry or cancel the removal request.

The **General Settings** window lets you change some of the general system-wide settings for your PingFederate servers.

About this task

The settings in this window are independent of one another.

Steps

- 1. Go to System # Server # General Settings.
- 2. Change any of the following settings:
 - To turn off automatic multi-connection error checking, select the Disable Automatic Connection Validation check box. This check box is not selected by default.
 - To override the verbosity of runtime transaction logging for all identity provider connections, select an option on the IdP Connection Transaction Logging Override menu. The default setting is Don't Override.
 - To override the verbosity of runtime transaction logging for all service provider connections, select an option on the SP Connection Transaction Logging Override menu. The default setting is **Don't Override**
 - To fine tune the caching interval for datastore validation, change the number of seconds in the Data-Store Validation Interval field. A value of 0 turns off the caching and validation tests are executed with each access. The default value is 300 seconds (five minutes). This setting applies to all datastores.



Note:

As you configure components on the administrative console, PingFederate performs connectivity tests against the applicable datastores. By default, PingFederate saves successful test results for five minutes. This design improves the performance of the administrative console by reducing the number of calls it makes to the target servers and the amount of time it takes to move from one configuration window to another.

• To let PingFederate use the value of a specific request header for the request ID instead of generating a request ID, enter the name of the request header in the Request Header for Correlation ID field. By default, the field is empty.



Note:

When a PingFederate endpoint receives HTTP requests, it records request IDs at the DEBUG level of the server log. When the Request Header for Correlation ID field specifies a request header, if a PingFederate endpoint receives a request that includes that header, and if the header's value contains 1 to 50 alphanumeric characters and hyphens, then PingFederate uses that value for the request ID. Otherwise, PingFederate uses a unique value that it generates for the request ID.

If the request causes PingFederate to send an LDAP call to PingDirectory, PingFederate includes the request ID in the LDAP call. For auditing and troubleshooting, you can correlate the request ID in the PingFederate log with the request ID in the PingDirectory log. For more information, see Correlating PingFederate events with PingDirectory LDAP activities on page 944.

3. Save the changes.

You can configure metadata settings, export metadata, or sign metadata XML files.

These menu items are located in the **Protocol Metadata** window.

Metadata settings

Configure metadata settings using the Pingfederate administrative console.

In the Metadata Settings window, you can configure the contact information to be included in your SAML metadata, and the metadata signing policy for metadata provided by the /pf/ federation metadata.ping federation metadata endpoint. You can also configure the validity of manual metadata exports and metadata provided by the metadata endpoint, and the frequency of automatic reloading of SAML metadata from partners.

Entering system information

Provide the contact information to be included in your SAML metadata.

Steps

- 1. Go to System # Protocol Metadata.
- 2. In the **Metadata Settings** window, on the **System Info** tab, enter the desired information.
- Click Next.



When editing an existing configuration, you can also click **Save** as soon as the administrative console offers the opportunity to do so.

Result

PingFederate includes the contact information in the manual metadata exports and metadata provided by the /pf/federation metadata.ping federation metadata endpoint.

Configuring metadata signing

Configure metadata signing using the PingFederate administrative console.

About this task

PingFederate generates publicly available metadata for partners through the federation metadata endpoint, /pf/federation metadata.ping. Although optional, signing the the metadata is recommended so that partners can verify the authenticity of the metadata.

Steps

- 1. Go to System # Protocol Metadata.
- 2. In the Metadata Settings window, on the Metadata Signing tab select a certificate from the Signing Certificate list.

If you have not yet created or imported your certificate into PingFederate, click Manage Certificates and use the **Certificate Management** configuration wizard to complete the task.

3. Optional: Select a signing algorithm from the list.

The default selection is RSA SHA256 or ECDSA SHA256 depending on the key algorithm of the chosen signing certificate. Make a different selection if you and your connection partner have agreed to use a stronger algorithm.

The public key of the metadata signing certificate is included as part of the metadata.

4. Click Next.



Tip:

When editing an existing configuration, you can also click Save as soon as the administrative console offers the opportunity to do so.

Configuring metadata lifetime

Configure the lifetime of metadata to optimize cache performance in your environment.

About this task

PingFederate provides metadata for SAML and WS-Federation connections and supports automatic update for SAML connections by reloading metadata URLs provided by the partners

Metadata publication

PingFederate includes expiration information in metadata. It indicates to partners whether they have reasonably up-to-date information about your server.

Metadata consumption

PingFederate supports automatic reloading of metadata by URL for SAML connections.

Steps

- 1. Go to System # Protocol Metadata.
- 2. In the Metadata Settings window, on the Metadata Lifetime tab, configure your metadata.

Option	Action
Adjust the validity of your metadata	Modify the Cache Duration field value, in minutes.
	The default value is 1440 or 1 day.
Adjust the frequency of automatic reloading of SAML metadata	Modify the Reload Delay field value, in minutes.
	The default value is 1440 or 1 day.

3. Click Next.



When editing an existing configuration, you can also click Save as soon as the administrative console offers you the opportunity to do so.

Reviewing metadata settings

Review your metadata settings before completing configuration.

Steps

- To amend your configuration, click the corresponding tab title and then follow the steps to complete the task.
- To keep your changes, click Save.
- To discard your changes, click Cancel.

Use the Pingfederate administrative console to export metadata.

The SAML standards define a metadata exchange schema for conveying XML-formatted information between two SAML entities. Metadata includes endpoint URLs, binding types, attributes, and securitypolicy information, which helps federation partners expedite their configurations.

In the **Metadata Export** window, you can export metadata to an XML file by selecting any SAML browser single sign-on (SSO) connection or specifying the desired information manually. The former is also available as a per-connection action item on the **Connections** window. The latter addresses the scenarios where you have not yet created a SAML connection, or you want to generate one SAML metadata XML file for multiple partners.

To export a SAML metadata file, select the role your PingFederate server plays, configure the export options and metadata signing policy, and then save the SAML metadata to an XML file.

For more information on exporting metadata for any SAML Browser SSO connection to an XML file, see Exporting connection-specific SAML metadata on page 860.

For more information on manually selecting specific information and exporting a metadata XML file, see Exporting selected SAML metadata on page 861.

Exporting connection-specific SAML metadata

You can export metadata for any SAML browser single sign-on (SSO) connection to an XML file.

About this task

This is useful in a situation where you have already created a SAML browser SSO connection to your partner and the partner prefers consuming SAML metadata by file.

Steps

- 1. Go to System # Protocol Metadata # Metadata Export.
- 2. On the **Metadata Role** tab, select the applicable role.
- 3. On the Metadata Mode tab, select the Use a connection for metadata generation option.

If the secondary HTTPS port is configured and you want to use it for the SOAP channel, select the Use the secondary port for SOAP channel check box.



Note:

If certificate-based authentication is configured for the SOAP channel, you must configure the pf.secondary.https.port property in the <pf install>/pingfederate/bin/ run.properties file and select this check box.

- 4. On the Connection Metadata tab, select the applicable SAML browser SSO connection from the list. Choose from:
 - Virtual Server ID

If the selected connection contains two or more virtual server IDs, you must select the virtual server ID that you want to use during the export. The protocol endpoints in the metadata file are specific to the selected virtual server ID. If you decide to update the virtual server ID at a later time, re-export the connection metadata for your partners

Virtual Host Name

If PingFederate is configured with one of more virtual server host names, you can select the applicable virtual host name from the list. If a selection is made, PingFederate use that virtual host name when generating the metadata file. If left blank, PingFederate uses its base URL in the metadata file. If you decide to update one or more virtual host names at a later time, re-export the connection metadata for your partners.

a. Select a certificate from the Signing Certificate list.

If you have not yet created or imported your certificate into PingFederate, click Manage Certificates and use the Certificate Management configuration wizard to complete the task.

- b. Optional: Select the related check boxes to include the public key information and the raw key in the signed XML file.
- c. Select a signing algorithm from the list.

The default selection is RSA SHA256 or ECDSA SHA256, depending on the key algorithm of the chosen signing certificate. Make a different selection if you and your connection partner have agreed to use a stronger algorithm.

- 6. On the Export & Summary tab, click Export to save the metadata XML file, then click Done.
- 7. Pass the metadata XML file to your partner.

Exporting selected SAML metadata

You can manually select the desired information and export a metadata XML file.

About this task

This type of export is useful for the following situations:

- You have not yet created a SAML browser single sign-on (SSO) connection to the partner but would like to help your partner with its configuration by including selected information in a metadata XML file.
- You want to export a SAML metadata with selected information, which can be passed to multiple partners to expedite their configurations.

Steps

- 1. Go to System # Protocol Metadata # Metadata Export.
- 2. On the **Metadata Role** tab, select the applicable role.
- 3. On the Metadata Mode tab, select the Select information to include in metadata manually option.

If the secondary HTTPS port is configured and you want to use it for the SOAP channel, select the Use the secondary port for SOAP channel check box.



Note:

If certificate-based authentication is configured for the SOAP channel, you must configure the pf.secondary.https.port property in the <pf install>/pingfederate/bin/ run.properties file and select this check box.

- 4. On the **Protocol** tab, select the desired version of the SAML protocol from the list.
- 5. On the Virtual Host Name tab, select the applicable virtual host name from the list.

Shown and applicable only if PingFederate is configured with one of more virtual server host names.

If a selection is made, PingFederate use that virtual host name when generating the metadata file. If left blank, PingFederate uses its base URL in the metadata file. If you decide to update one or more virtual host names at a later time, re-export the connection metadata for your partners.

6. Optional: On the Attribute Contract tab, you can perform the following actions.

Action	Description
Add	Add an attribute contract by entering the contract's name and clicking Add .

Action	Description
Edit	Modify an existing attribute contract by clicking Edit . To save your change, click Update . To cancel your change, click Cancel .
Delete	Delete an existing attribute contract by clicking Delete .

7. Optional: On the **Signing Key** tab, if you want to include a public key that this system uses for digital signatures, select an available key from the **Digital Signature Keys/Certs** list.

If you have not yet created or imported a digital signature key to PingFederate, click **Manage Certificates** and use the **Digital Signature Settings** wizard to complete the task.

- 8. Optional: On the **Metadata Signing** tab, select a certificate to use for signing the metadata XML file.
 - a. Select a certificate from the Signing Certificate list.
 - If you have not yet created or imported your certificate into PingFederate, click **Manage**Certificates and use the Certificate Management configuration wizard to complete the task.
 - b. Optional: Select the related check boxes to include the public key information and the raw key in the signed XML file.
 - c. Select a signing algorithm from the list.
 - The default selection is **RSA SHA256** or **ECDSA SHA256**, depending on the key algorithm of the chosen signing certificate. Make a different selection if you and your connection partner have agreed to use a stronger algorithm.
- 9. Optional: On the **XML Encryption Certificate** tab, select the certificate that your partner can use to encrypt XML content.
 - Applicable only when you have selected **SAML 2.0** on the **Protocol** tab.
 - If you have not created or imported your certificate into PingFederate, click **Manage Certificates** and use the **Certificate Management** configuration wizard to complete the task.
- 10. On the **Export & Summary** tab, click **Export** to save the metadata XML file, then click **Done**.
- 11. Pass the metadata XML file to your partner or partners.

File signing

Applying a digital signature to an XML file assures the authenticity and integrity of the original source.

If you have previously exported an unsigned metadata XML file, you can create a new signed metadata XML file based on the unsigned metadata in the **File Signing** window.

Signing XML files

Use the Pingfederate administrative console to sign XML files.

Steps

- 1. Go to System # Protocol Metadata # File Signing.
- 2. On the **Select Metadata File** tab, choose your metadata file.
- 3. On the **Digital Signature Settings** tab, select a signing certificate from the list.

If you have not yet created or imported your certificate into PingFederate, click **Manage Certificates** and use the **Certificate Management** configuration wizard to complete the task.

- a. Clear the related check boxes to exclude the public key information and the raw key from the signed XML file.
- b. Select a signing algorithm from the list.
 - The default selection is **RSA SHA256** or **ECDSA SHA256**, depending on the key algorithm of the chosen signing certificate. Make a different selection if you and your connection partner have agreed to use a stronger algorithm.
- 4. On the Export & Summary tab, click Export to save the digitally signed file.

Monitoring and notifications

You can configure PingFederate to process or send SNMP requests from or to your SNMP networkmanagement server.

You can also set up notifications for licensing events, certificate events, and SAML metadata update events. These menu items are located in **System # Monitoring & Notifications**.

Runtime notifications

You can configure PingFederate to generate notification messages for various events.

Complete this configuration on the **Runtime Notifications** window.

Licensing events

Depending on your licensing agreement, your PingFederate license might have an expiration date. You can configure PingFederate to generate notification messages when your license is about to expire. This option does not appear when you have a perpetual license.

Certificate events

When enabled, PingFederate can generate notification messages when a certificate is about to expire or has expired. If you have also configured PingFederate to rotate self-signed certificates, PingFederate can generate notification messages after generating a certificate and after its activation.

SAML metadata update events

If you have enabled automatic reloading of partner metadata in any SAML browser single signon (SSO) connections, you can configure PingFederate to generate notification messages when it detects changes after pulling the metadata from the partners.



Note:

In a clustered PingFederate environment, notification messages for licensing and certificate events are generated by only one of the nodes. If that node leaves the cluster, planned or not, another node picks up this task automatically.

Regardless of runtime notification settings, PingFederate always records license expiration events, certificate events, and SAML metadata update events in the server log. For more information, see PingFederate log files on page 761.

Configuring runtime notifications

Use the PingFederate administrative console to configure and manage runtime notifications.

Steps

1. Go to System # Monitoring & Notifications # Runtime Notifications .

- 2. Optional: Configure licensing events.
 - Select the Notification for Servers Licensing Events check box.



Note:

This check box appears only if your PingFederate license has an expiration date.

- b. Enter the email address of the intended recipient.
- c. Select a notification publisher instance from the list.

If you have not yet configured the desired notification publisher instance, click Manage Notification Publishers.

- 3. Optional: Configure certificate events.
 - a. Select the Notification for Certificate Events check box.
 - b. Enter the email address of the intended recipient.
 - c. Optional: Enter the initial warning time period (in days) in the **Initial Warning Event** field.
 - d. Optional: Enter the final warning time period (in days) in the Final Warning Event field.
 - Select a notification publisher instance from the list.

If you have not yet configured the desired notification publisher instance, click Manage Notification Publishers.

- 4. Optional: Configure SAML metadata update events.
 - a. Select the Notification for SAML Metadata Update Events check box.
 - b. Enter the email address of the intended recipient.
 - c. Select a notification publisher instance from the list.

If you have not yet configured the desired notification publisher instance, click Manage Notification Publishers.

Runtime reporting

PingFederate supports runtime monitoring and reporting through the Simple Network Management Protocol (SNMP), a standard used by network-management consoles to monitor network and server activity across an enterprise.

PingFederate also supports runtime monitoring and reporting through Java Management Extensions (JMX).



You can use HTTP requests at any time to verify the status of the PingFederate server. For more information, see Customizing the heartbeat message on page 835.

You can also supplement monitoring information by applying third-party analysis and reporting tools to the security audit log, in which PingFederate records fine-grain details, including response times and event types, for all server transactions. For more information, see Security audit logging on page 767.

Configuring SNMP monitoring

Configure SNMP monitoring to get an overview of the network performance.

About this task

The SNMP management information base (MIB) defines network data available for SNMP monitoring. The MIB file is located in the <pf install>/pingfederate/SNMP directory. The MIB describes the object identifiers that PingFederate uses to communicate information through SNMP. These identifiers are globally unique and managed by the Internet Assigned Numbers Authority (IANA).

SNMP supports Gets and Traps.

A Get is a request for status information sent by a network-management console to an SNMP agent. Embedded within each PingFederate server is an SNMP agent that brokers the communication between the management console and the PingFederate runtime engine.

PingFederate responds to two single sign-on (SSO) and single logout (SLO) types of Get requests:

- The total number of transactions that the server instance has processed since installation.
- The total number of failed transactions that the server instance has encountered since installation

Because PingFederate is built within an existing Jetty framework, Gets include a variety of server information available through Jetty-standard Managed Beans (MBeans). A detailed list of this information is provided in the MIB file in the <pf install>/pingfederate/SNMP directory. For more information about MBeans, see Runtime monitoring using JMX on page 865.



If the host operating system does not allow the SNMP agent to bind to privileged ports (those below 1024), consult your operating system's documentation on how to work around this limitation or change the default port 161 to a port above 1023.

Traps

A Trap is a spontaneous communication from an agent to a network-management console. PingFederate generates a Trap at regular intervals, the server "heartbeat." Each Trap contains the amount of time the server instance has been running since its most recent start-up.

Steps

- 1. Go to System # Monitoring & Notifications
- 2. Optional: In the **Runtime Reporting** window, enable Gets.
 - a. Select the **Respond to Get Requests** check box.
 - b. Modify the Local Agent Port and the Community Name field values as needed.
- 3. Optional: In the **Runtime Reporting** window, enable Traps.
 - a. Select the **Generate Traps** check box.
 - b. Provide the required information for your network-management console and modify the "Heartbeat" Interval field value as needed.
 - c. Click Test SNMP Configuration to send a single Trap to your network-management console
 - d. Verify the result.
- 4. To keep your changes, click Save.
- 5. To discard your changes, click **Cancel**.

Runtime monitoring using JMX

Similar to SNMP, Java Management Extensions (JMX) technology represents a Java-centric approach to application management and monitoring.

JMX exposes instrumented code in the form of MBeans. Application management systems that support JMX technology, such as JConsole, can request runtime information from the PingFederate JMX server.



Important:

Authentication is required for JMX-client access to PingFederate runtime data. For more information, see Configuring service authentication on page 636.

SSO and SLO monitoring

For SSO/SLO transaction processing, PingFederate provides these MBeans:

- pingfederate:type=TOTAL FAILED TRANSACTIONS
- pingfederate:type=TOTAL TRANSACTIONS

Each type contains a single attribute, Count, which reports the same information as an SNMP Get.

Sample Jetty metrics

The following table describes examples of Jetty MBean metrics, available through JMX, that administrators might find useful to supplement information provided through the PingFederate-specific MBeans.

Attributes
connections – The total number of TCP connections accepted by the server.
connectionsDuration* – How long connections are kept open. Maximum, mean, standard deviation, and total accumulated time are available.
connectionsOpen - The current number of open connections. Maximum is also available (connectionsOpenMax).
mequests - Total number of requests received.
requestsActive — Number of requests currently being processed. Max is also available.
$\verb requestTime-Request \ duration. \ Maximum, \ mean, \ standard \ deviation, \ and \ total \ accumulated \ time \ are \ available.$
responses1xx, responses2xx, responses3xx, – Total number of requests that returned HTTP status codes of 1xx, 2xx, 3xx, etc.
idleThreads - Number of idle threads currently available.
threads – Number of threads currently running, including both idle and active.
minThreads - Minimum number of threads in the pool.
maxThreads — Maximum number of threads in the pool.
${\tt lowOnThreads-A\ boolean\ flag\ indicating\ whether\ the\ pool} \\ is\ running\ low\ on\ threads. \\$
Various attributes measuring CPU usage and memory.

Advanced JMX configuration

PingFederate uses port 1099 for its JMX server. To change the port or other Java Message Service (JMS) configuration items, if needed, modify the jmx-remote-config.xml configuration file in the <pf install>/pingfederate/server/default/conf directory.



When connecting to the JMX service using SSL, the default, ensure that the client trusts the PingFederate SSL server certificate presented. For more information, see Manage SSL server certificates on page 601.

Datastores

Datastores represent external systems where user attributes and other data are stored. Once defined, you can configure PingFederate to retrieve user attributes from datastores for contract fulfillment and token authorization in various use cases.

To manage datastores, go to System # Data & Credential Stores # Data Stores.

- To create a new datastore, click Add New Data Store and then follow the configuration wizard to complete the task.
- To modify an existing datastore, select the datastore and then follow the configuration wizard to complete the task.
- To review usage of an existing datastore, click Check Usage under Action.
- To remove an existing datastore or cancel the removal request, click Delete or Undelete under Action.



Note:

You can only remove datastores that are not currently in use.

 To fine-tune the caching interval for datastore validation, update the Data-Store Validation Interval field value to the desired amount of time in seconds.



Note:

As you configure various components on the administrative console, PingFederate performs connectivity tests against the applicable datastores. By default, PingFederate stores successful test results for five minutes. This design improves the performance of the administrative console by reducing the number of calls it makes to the target servers and the amount of time it takes to move from one configuration window to another.



Note:

The default value is 300 seconds (five minutes). A value of 0 turns off the caching and validation tests are executed with each access. This setting applies to all datastores.

Adding a new datastore

Create and configure a datastore.

Steps

- 1. Go to System # Data & Credential Stores # Data Stores.
- 2. Click Add New Data Store.
- 3. Enter a name for the datastore.

- 4. From the **Type** list, select the type of datastore. Available types are limited to the ones currently installed on your server.
- 5. Optional: To mask attribute values returned from this datastore in PingFederate logs, select the Mask Values in Log check box.
- 6. Click Next.

Configuring a JDBC connection

Provide the required information to establish a Java Database Connectivity (JDBC) connection to your database server.

About this task



PingFederate was tested with vendor-specific JDBC drivers. For more information, see *Database driver* information on page 115. To obtain the database driver JAR file, contact your database vendor. Install the database driver file to the <pf install>/pingfederate/server/default/lib directory then restart the server.

Steps

- 1. Go to System # Data & Credential Stores # Data Stores.
- 2. On the Data Stores window, click Add New Data Store.
- 3. On the **Data Store Type** tab, type a name for the datastore.
- 4. From the **Type** list, select **Database (JDBC)**. Click **Next**.
- 5. Optional: To mask attribute values returned from this datastore in PingFederate logs, select the Mask Values in Log check box.
- 6. Click Next.
- 7. In the Database Config window, configure your JDBC connection. Information about each field is provided in the following table.

Field	Description
JDBC URL	The location of the database server and the database. The structure of the JDBC URL varies depending on the vendor. You can add multiple JDBC URLs. You can also specify which node is the default by clicking Set as Default under Action .
	Tip:
	For Oracle MySQL, to enable automatic reconnection attempts when the connection is not available at runtime, enter a SQL statement in the Validate Connection SQL field and add the following query string to the JDBC URL:
	?autoReconnect=true
Tags	Tags are defined in the node.tags property in the <pf_install>/ pingfederate/bin/run.properties file. For a description of the node.tags property, see Deploying cluster servers on page 201.</pf_install>
	In PingFederate deployments that are regional, you can enter one or more tags for a JDBC URL, which specifies with which datastore that particular PingFederate node should communicate. If none of the tags

Field	Description
	match what is defined for the node.tags property, the default node is used.
	The following rules apply to tags:
	 You must separate multiple tags specified for one node with spaces. You can't use a tag more than once per datastore. Tags are optional. If needed, you can configure a non-default node without tags. Doing this is useful if you are not yet ready to tag the node, or if you are still in the planning stage but want to enter the address for the node now.
Driver Class	The name of the driver class used to communicate with the source database. The driver class name should be supplied by the database software vendor in a JAR file.
Username	The name that identifies the user when connecting to the database.
Password	The password needed to access the database.
Validate Connection SQL	A simple SQL statement used by the PingFederate runtime server to verify that the database connection is still active and to reconnect if needed.
(Optional but recommended)	If a SQL statement is not provided here, PingFederate might not reconnect to the database if the connection is broken.
	Important:
	Ensure that the SQL statement is valid for your database. For example:
	 SELECT 1 from dual (for Oracle Database or Oracle MySQL) SELECT getdate() (for Microsoft SQL Server) SELECT 1 (for PostgreSQL)
	Tip:
	To use this feature for Oracle MySQL, you must also add the ? autoReconnect=true query parameter to the JDBC URL.
Mask Values in Log	Determines whether all attribute values returned through this datastore should be masked in PingFederate logs.
	Applicable only when editing an existing datastore.
Allow Multi-Value Attributes	When selected, indicates that the JDBC datastore can select more than one record from a column and return the results as a multivalued attribute. Otherwise, a query returns only the first value in the column.

8. Click **Test Connection** to determine whether the administrative node can communicate with the specified datastore.



Datastore validation is no longer enabled during configuration. This feature lets you configure datastores without requiring a successful connection between the administrative node and the datastore. You can also save the datastore even if the connection is not currently successful.

- 9. Click **Advanced** to configure additional settings.
 - a. On the Advanced Database Options window, click Apply Defaults to view or restore default values.



The default values are conservative based on the server thread pool settings configured in the <pf install>/pingfederate/etc/jetty-runtime.xml file. If any changes are made to thread pooling, we recommend updating settings as outlined in the next step.

b. Configure advanced settings.

For more information about each field, see the following table.

Field	Description
Minimum Pool Size	The smallest number of database connections that can remain in the pool for the given datastore. A minimum value of 0 means that the minimum number of connections in the pool is zero.
	Note:
	For optimal performance, the value for this setting should equal 50% of the maxThreads value in the Jetty server configuration (<i>Configuring connection pools to datastores</i> on page 993).
	Note that PingFederate does not establish the connection pool for the given datastore until it receives a request that requires one or more attributes from that datastore.
	The default value (after clicking on Apply Defaults) is 10.
Maximum Pool Size	The largest number of database connections that can remain in the pool for the given datastore.
	Note:
	For optimal performance, the value for this setting should equal 75% to 100% of the maxThreads value in the Jetty server configuration (see Configuring connection pools to datastores on page 993).
	The default value (after clicking on Apply Defaults) is 100.
Blocking Timeout (ms)	The amount of time a request waits to get a connection from the connection pool before it fails. A value of −1 means that a request waits indefinitely for the connection pool to return a connection.
	The default value (after clicking on Apply Defaults) is 5000.

Field	Description
Idle Timeout (min)	The length of time the connections can sit idle in the pool before it closes them. A value of -1 means that the connection pool does not close its connections (once established).
	Note that PingFederate maintains the minimum connection pool for the given datastore once the pool is established.
	The default value (after clicking on Apply Defaults) is 5.

10. Click **Save** to save your configuration.

Configuring an LDAP connection

Provide the required information to establish an LDAP connection to your directory server.

Steps

- 1. Go to System # Data & Credential Stores # Data Stores.
- 2. On the Data Stores window, click Add New Data Store.
- 3. On the **Data Store Type** tab, type a name for the datastore.
- 4. From the Type list, select Directory (LDAP).
- 5. Optional: To mask attribute values returned from this datastore in PingFederate logs, select the **Mask Values in Log** check box.
- 6. Click Next.
- 7. On the LDAP Configuration tab, configure your LDAP connection as described in the following table.

Field	Description
Data Store Name	The name of the datastore.
	Applicable only when editing an existing datastore.
Hostname(s) (Required)	The network address of the directory server, either an IP address, a host name, or a fully qualified domain name. The entry might include a port number; for example, 10.10.10.101:1389. For failover, enter multiple directory servers, each separated by a space. In addition to network error conditions, PingFederate also fails over to the next server if the current server returns an LDAP system error.
	Note: If multiple directory servers are specified, each server must be accessible by using the same user distinguished name (DN) and password (unless the Bind Anonymously check box is selected).
	You can add multiple hostnames. You can also specify which node is the default by clicking Set as Default under Action .
	PingFederate can also leverage DNS service records to locate the directory server (when the Use DNS SRV Record check box is selected), in which case the value of this field must be a single domain; for example, example.com.

Field	Description
Tags	Tags are defined in the node.tags property in the <pf_install>/ pingfederate/bin/run.properties file. See Deploying cluster servers on page 201 for a description of the node.tags property.</pf_install>
	In regional PingFederate deployments, you can enter one or more tags for a host name, which specify with which datastore that particular PingFederate node should communicate. If none of the tags match what is defined for the node.tags property, the default node is used.
	The following rules apply to tags:
	 You must separate multiple tags specified for one node with spaces. You cannot use a tag more than once per datastore. Tags are optional. If needed, you can configure a non-default node without tags. This is useful if you are not yet ready to tag the node, or if you are still in the planning stage but want to enter the address for the
Use LDAPS	node now. When selected, PingFederate connects to the directory server using LDAPS. This selection applies equally to all servers specified in the Hostname(s) field.
	Important:
	We recommend securing all LDAP connections by using LDAPS.
	Note:
	To enable the password changes, password reset, or account unlock features in the HTML Form Adapter against Microsoft Active Directory, you must secure the connection to your directory server using LDAPS; Microsoft Active Directory requires this level of security to allow password changes.
	This check box is not selected by default.
Use DNS SRV Record	Used in conjunction with the domain information defined in the Hostname(s) field and the preference of LDAP or LDAPS, PingFederate uses DNS SRV records to locate the directory server when this check box is selected. You can fine-tune the TTL value and the record prefixes on the Advanced LDAP Options window.
	Note:
	When the DNS returns multiple SRV records, PingFederate uses the record with the lowest-numbered priority value and fails over to the record with the next lowest priority value. If multiple records share the same priority value, PingFederate uses the records with the highest-numbered weight value.
	PingFederate repeats this exercise until it establishes a connection or fails to connect to any directory server after taking all records into consideration.

This check box is not selected by default.

Field Description **Follow LDAP** Select this check box to let the datastore follow LDAP referrals on Microsoft Active Directory, Oracle Unified Directory, or Oracle Directory Server. Referrals Note: **LDAP Type** (Required) Tip: Adapter instance is used. **Bind Anonymously** Z Tip:

PingFederate always follows LDAP referrals from PingDirectory based on the recommended PingDirectory configuration.

If you are using this datastore for outbound provisioning and your directory server is PingDirectory, Microsoft Active Directory, Oracle Unified Directory, or Oracle Directory Server, select the applicable type from the list, such that PingFederate can pre-populate many provisioning settings on Outbound Provisioning # Channel # Source Settings.

If your directory server is not PingDirectory, Microsoft Active Directory, Oracle Unified Directory, or Oracle Directory Server, you can define a custom LDAP Type to streamline the outbound provisioning configuration.

The LDAP type is also used to enable password-change messaging between Microsoft Active Directory and PingFederate when an HTML Form

Select this check box if your directory server supports anonymous binding and if no credentials are needed to access the directory server. When selected, user DN and password are not required.



For inbound provisioning, because PingFederate needs to manage local user records, your directory server might require a specific service account to handle the communication between PingFederate and the target directory server. If you choose an anonymous binding, ensure that this access level provides permission to search the directory for user-account information.

This check box is not selected by default.

8. Click **Test Connection** to determine whether the administrative node can communicate with the specified datastore.



Note:

Datastore validation is no longer enabled during configuration. This feature lets you configure datastores without requiring a successful connection between the administrative node and the datastore. You can also save the datastore even if the connection is not currently successful.

- 9. Optional: Click Advanced. If you choose an anonymous binding, configure additional settings in the Advanced LDAP Options window.
- 10. Click Save.

Setting advanced LDAP options

PingFederate enables you to customize the default settings of both the search pool and the bind pool for each LDAP datastore.

About this task

PingFederate maintains a search pool and a bind pool for each LDAP datastore for optimal performance. The search pool is for LDAP directory searches. The bind pool is for LDAP bind authentication purposes. Use the Advanced LDAP Options window to change default pool settings. These settings are applicable to both the search pool and the bind pool.

When configuring PingFederate to locate the directory server based on DNS SRV record, you can finetune the TTL value and the SRV record prefixes.

1. In the Advanced LDAP Options window, click Apply Defaults to view or restore default values.



The default values are conservative based on the server thread pool settings configured in the $<\!pf_install\!>\!$ /pingfederate/etc/jetty-runtime.xml file. If any changes are made to thread pooling, update the settings as outlined in the following step.

2. Configure advanced settings. For more information about each field, see the following table.

Field	Description
Test Connection on Borrow	Indicates whether to validate objects before they are borrowed from the pool.
	This check box is not selected by default.
Test Connection on Return	Indicates whether to validate objects before they return to the pool.
	This check box is not selected by default.
Create New Connection If Necessary	Indicates whether you can create temporary connections when the Maximum Connections threshold is reached. Temporary connections are managed automatically.
	Note:
	If disabled, when the Maximum Connections value is reached, subsequent requests relying on this LDAP datastore instance might fail.
	This check box is selected by defaul
Verify LDAPS Hostname	Indicates whether to verify that the host name of the directory server matches the subject (CN) or one of the subject alternative names (SANs) from the certificate.
	Important:
	Verify the LDAPS host name for all LDAPS connections.
	This check box is selected by default.
Minimum Connections	The smallest number of connections that can remain in each pool. A minimum value of 1 creates two connections, one
(Required)	connection in the search pool and one connection in the bind pool. The default value is 10.
	Note:
	For optimal performance, the value for this setting should equal 50% of the maxThreads value in the Jetty server configuration. For more information see <i>Configuring connection pools to datastores</i> on page 993.

Field	Description
	Note: PingFederate does not establish the connection pool for the
	given datastore until it receives a request that requires one or more attributes from that datastore.
Maximum Connections	The largest number of active connections that can remain in each pool (not including the temporary connections that are
(Required)	managed automatically when the Create New Connection If Necessary check box is selected). The value must exceed or equal the Minimum Connections value.
	Note:
	For optimal performance, the value for this setting should equal 75% to 100% of the maxThreads value in the Jetty server configuration. For more information, see <i>Configuring connection pools to datastores</i> on page 993.
	The default value is 100.
Maximum Wait (Milli)	The maximum number of milliseconds the pool waits for an available connection when trying to obtain a connection from
(Required)	the pool. A value of -1 causes the pool not to wait at all and to either create a new connection or produce an error (when no connections are available).
	The default value is −1.
Time Between Eviction (Milli) (Required)	The number of milliseconds between periodic background health checks against the available connections in this pool. A value of -1 disables the evictor.
	The default value is 60000.
Read Timeout (Milli) (Required)	The maximum number of milliseconds a connection waits for a response to return before producing an error. A value of -1 causes the connection to wait indefinitely.
	The default value is 3000.
Connection Timeout (Milli) (Required)	The maximum number of milliseconds that a connection attempt can continue before returning an error. A value of -1 causes the pool to wait indefinitely.
	The default value is 3000.
DNS TTL (Milli) (Required)	The amount of time in milliseconds that a previously obtained DNS SRV record remains valid. When this threshold is reached, PingFederate contacts the DNS for a new SRV record to locate the directory server.
	The default value is 60000.
LDAP DNS SRV Record prefix	The prefix that PingFederate uses in its DNS queries for SRV records to locate an LDAP-capable directory server.

Field	Description
(Required)	The default value is _ldaptcp.
LDAPS DNS SRV Record prefix	The prefix that PingFederate uses in its DNS queries for SRV records to locate an LDAPS-capable directory server.
(Required)	The default value is _ldapstcp.

- 3. Optional: Click Next to specify LDAP binary attributes on the LDAP Binary Attributes tab.
- 4. Click Save.

Specifying LDAP binary attributes

PingFederate allows you to specify attributes where you must handle such attribute values as binary data for use in attribute contract fulfillment.

About this task

You cannot use binary data in an assertion. You must apply and handle encoding on a per-connection basis. When binary attributes are selected for attribute mapping, the administrative console prompts you to select an encoding type for each binary attribute.

Steps

- 1. On the LDAP Binary Attributes window, add, edit, or remove binary attributes.
- 2. Click **Save** to keep your configuration.

Proxied authorization

When connecting to PingDirectory or Oracle Directory Server, configure proxied authorization for the service account on the directory server if you intend to enable self-service password reset in any HTML Form Adapter instances that use this datastore.

By configuring proxied authorization for the service account on the directory server, users are not allowed to reset their passwords if their accounts are disabled or if they were not granted permission to change their passwords.

For information on configuring proxied authorization for service accounts, see the following table.

Directory server	Reference
PingDirectory	See Working with Proxied Authorization in the PingDirectory Administration Guide.
Oracle Directory Server	Go to Oracle's Oracle Fusion Middleware Deployment Planning Guide and search for "Proxy Authorization".
Oracle Unified Directory	Go to Oracle's online guide https://docs.oracle.com/cd/E52734_01/oud/OUDAG/toc.htm Fusion Middleware Administering Oracle Unified Directory and search for "proxied authorization control" in its glossary .



Note:

Microsoft Active Directory does not support proxied authorization. See https://docs.microsoft.com/en-us/openspecs/windows_protocols/ms-adts/faf0b8c6-8c59-439f-ac62-dc4c078ed715?redirectedfrom=MSDN.

For general information about proxied authorization, see RFC4370.

Configuring the password validation details request control ACI

When connecting to PingDirectory, configure the password validation details request control Access Control Instruction (ACI) to provide user-friendly messages when users fail to change or reset their

passwords through the self-service account management capabilities in any HTML Form Adapter instances that use the datastore.

About this task

For self-service password management, where the user knows the current password and wants to update it, the service account of the datastore must have the password validation details request control ACI. For self-service account recovery, where the user wants to define a new password after forgetting the current password, the user account needs the same ACI.

Steps

1. Create LDIF files to capture the following ACI information.

OID

```
1.3.6.1.4.1.30221.2.5.40
```

Name

Password Validation Details Requerst Control

Permission

read

The following examples show the example file contents for change password and password reset.

Example: aci toSvcAccount forChangePassword.ldif

```
# ACI to service account for change password
dn: uid=ssoDataStore, ou=ServiceAccounts, dc=example, dc=local
changetype: modify
add: aci
aci: (targetcontrol="1.3.6.1.4.1.30221.2.5.40") (version
 3.0; acl "Access to the Password Validation Details
Request Control"; allow (read) userdn="ldap:///
uid=ssoDataStore,ou=ServiceAccounts,dc=example,dc=local";)
```

Example: aci toUsrAccount forPasswordReset.ldif

```
# ACI to a user account for password reset
dn: uid=user.7,ou=People,dc=example,dc=local
changetype: modify
add: aci
aci: (targetcontrol="1.3.6.1.4.1.30221.2.5.40") (version 3.0; acl "Access
 to the Password Validation Details Request Control"; allow (read)
 userdn="ldap:///uid=user.7,ou=People,,dc=example,dc=local";)
```



Note:

For demonstration purposes, this sample LDIF file only targets one user. You can use other LDIF syntax to widen its coverage to include multiple users.

2. Use the ldapmodify command to configure the required ACI.

Example:

```
$ ldapmodify -f <path>/aci toSvcAccount forChangePassword.ldif
-h <host name>
-p <LDAP port>
-D <LDAP bind username>
```

```
-w <LDAP bind password>
$ ldapmodify -f <path>/aci toUsrAccount forPasswordReset.ldif
-h <host name>
-p <LDAP port>
-D <LDAP bind username>
-w <LDAP bind password>
```



Note:

Line breaks are inserted for readability only.

Defining a custom LDAP type for outbound provisioning

If you are using outbound provisioning and your directory server is not PingDirectory, Microsoft Active Directory, Oracle Unified Directory, or Oracle Directory Server, you can define a custom LDAP type for PingFederate to use to streamline the provisioning configuration.

Steps

- 1. Copy and rename <pf install>/pingfederate/server/default/conf/template/ldaptemplates/sample.template.txt file.
- 2. Change the template.name property value in the new template file. This property value appears in the LDAP Type list on the LDAP Configuration window when you save the template.
- 3. Modify other property values in the file to match the corresponding configuration of your directory server.
 - These properties correspond to the fields shown on **Outbound Provisioning # Channel # Source Settings**. They help the provisioner determine when user records are added, changed, or removed.
- 4. Save the new template file.
 - For a clustered PingFederate environment, perform these steps on the console node. No changes or restart of the PingFederate service is required on any nodes.

Next steps

After you have configured the LDAP type, you can create a new LDAP datastore using the newly defined LDAP type. To streamline outbound provisioning configuration, select the LDAP data store that uses the newly-defined LDAP type in the **Source** window.

Configuring other types of datastores

Besides connecting to a directory server using LDAP or a database server using Java Database Connection (JDBC), PingFederate can connect to other types of datastores, such as REST API-enabled data sources that return user attributes in JavaScript Object Notation (JSON). See the following topics for configuration steps:

- Configuring a REST API datastore on page 879
- Configuring a custom datastore on page 883

Configuring a REST API datastore

To retrieve attribute data from a JSON-based REST API, you must first create a REST API datastore.

Steps

- 1. Go to System # Data & Credential Stores # Data Stores.
- 2. On the **Data Stores** window, click **Add New Data Store**.
- 3. On the **Data Store Type** tab, type a name for the datastore.

- 4. From the **Type** list, select **Rest API**.
- 5. Optional: To mask attribute values returned from this datastore in PingFederate logs, select the Mask Values in Log check box.
- 6. Click Next.
- 7. On the Configure Data Store Instance tab, click Add a new row to 'Base URLS and Tags'.
 - a. Enter the Base URL of the data source offering REST API access to its data. You can enter multiple base URLs.
 - b. Optional: Enter one or more tags per base URL.

Tags are defined in the node.tags property in the <pf install>/pingfederate/bin/ run.properties file. For a description of the node.tags property, see Deploying cluster servers on page 201.

In PingFederate deployments that are regional, you can enter one or more tags for a Base URL, which specifies the PingFederate node the datastore should communicate with. If none of the tags match what is defined for the node.tags property, the default node is used.

The following rules apply to tags:

- You must separate multiple tags specified for one node with spaces.
- Tags must unique per base datastore.
- You cannot use a tag more than once per datastore.
- Tags are optional. If needed, you can configure a non-default node without tags. Doing this is useful if you are not yet ready to tag the node, or if you are still in the planning stage but want to enter the address for the node now.
- c. Click Update under Action.
- d. Select Set as Default under Action beside the Base URL and Tags that you want to use as the default. The first Base URL and Tags configured is set as the default automatically.



Note:

If the data source exposes multiple paths or requires specific query parameters to retrieve user records, enter the base URL here and then specify the path and query parameters in the attribute source configuration.

For more information, see Specifying a resource path for a REST API datastore on page 369.

8. If the data source requires specific HTTP request headers, click Add a new row to 'HTTP Request Headers'.



If configured, PingFederate includes the configured HTTP request headers and their values when contacting the data source.

- a. Enter the applicable name and value under Header Name and Header Value.
- b. Click **Update** under **Action**.

Repeat these steps to define additional HTTP request headers and their values.

9. Click Add a new row to 'Attributes' to define local attribute names and map them to the data returned by the data source.

Map each attribute to a path representing an attribute in the JSON response. This path follows the syntax defined in the JavaScript Object Notation (JSON) Pointer specification at tools.ietf.org/html/ rfc6901.

You must define at least one attribute.

- a. Enter the Local Attribute name and JSON Response Attribute Path.
- b. Click **Update** under **Action**.

Repeat these steps to define additional attributes.



Define only the attributes required by other configuration items, such as contract fulfillment or token authorization. Provide meaningful attribute names so that you can easily recognize them at a later time.

10. Select one of the following authentication methods.

Choose from:

None

PingFederate makes unauthenticated REST API requests to the data source. No credential information is required. This is the default setting.

Basic Authentication

PingFederate authenticates via the HTTP Basic authentication scheme. Enter the required credentials in the Username and Password fields.

OAuth 2.0 Bearer Token

PingFederate authenticates by presenting an OAuth 2.0 access token.

In this scenario, PingFederate is an OAuth client, specifically a client that uses the client credential grant type to obtain access token from an authorization server and presents the access token to the data source for authentication.

Enter the client credentials in the Client ID and Client Secret fields. Then enter the token endpoint URL at the authorization server and the applicable scope (or scopes) in the OAuth Token Endpoint and OAuth Scope fields.

11. If PingFederate should mask attribute values returned through this datastore in its log, select the Mask Values in Log check box.

This check box is visible only when editing an existing datastore and is not checked by default.

12. Optional: Click **Show Advanced Fields** to configure additional settings.

For more information, see the following table.

Field	Description
Enable HTTPS Hostname Verification	Indicates whether to verify that the hostname of the data source matches the subject (CN) or one of the subject alternative names (SANs) from the certificate.
	Important:
	We recommend to verify hostname for all connections.
	This check box is selected by default.

Field	Description
Read Timeout (MS)	Defines the socket timeout in milliseconds.
	Enter 0 to set an infinite timeout.
	Enter a negative integer to use the default value set by the operating system.
	The default value is 10000 in milliseconds, which is 10 seconds.
Connection Timeout	Determines the timeout in milliseconds until a connection is established.
(MS)	Enter 0 to set an infinite timeout.
	Enter -1 to use the default value set by the operating system.
	The default value is 10000 in milliseconds, which is 10 seconds.
Max Payload Size (KB)	Defines the maximum allowed size in kilobytes (KB) of the returned JSON response payload.
	Enter 0 to configure an unrestricted payload size.
	The default value is 1024 in KB.
Retry Request	Determines whether to retry a user data retrieval request if the data source returns an HTTP status code found in the Retry Error Codes .
	This check box is selected by default.
Maximum Retries Limit	Defines the maximum number of retry attempts if the data source returns an HTTP status code found in the Retry Error Codes .
	The default value is 5.
Retry Error Codes	Enter a comma-separated list of HTTP status codes, for which if received from the data source, PingFederate might retry the request.
	For example, you can enter 429 for "Too Many Request" or 503 for "Service Unavailable".
	The default value is 429.
Test Connection URL	Determines the URL to which PingFederate sends GET requests to test the datastore connection on the Actions tab.
	When not specified (the default), PingFederate sends GET requests to the base URL of the datastore.

- 13. On the **Actions** tab, verify the datastore configuration.
 - a. Click **Test Connection** to test the connectivity between PingFederate and the data source.
 Result:

The administrative console displays the results returned by the data source. The PingFederate server log may contain additional messages as well.

- b. Review the results.
- c. Optional: Click **Reset** and repeat the test.
- 14. On the **Summary** tab, review your configuration, amend as needed, click **Save** to keep your configuration or click **Cancel** to discard it.

You have two use cases that can leverage user attributes obtained through REST APIs. The data source returns user records in JSON

```
"uid": "asmith",
"office": {
 "city": "Denver",
 "state": "CO",
 "zipCode": 80202
"telephoneNumbers": [
  "+1 303-555-1234",
  "+1 303-555-5678"
"department": "Engineering"
```

The first use case requires the user's department, while the second use case requires the first telephone number and the ZIP code.

To address both use cases, create a REST API datastore with the following attributes.

Local Attribute	JSON Response Attribute Path
Dept	/department
Telephone	/telephoneNumbers/0
Zip	/office/zipCode

Once set up, you can fulfill various contracts or configure issuance criteria based on the attribute data from the data source.

Configuring a custom datastore

You can configure your own custom datastore instance to perform specified actions.

About this task

Developers can use the PingFederate SDK to create specific drivers for non-Java Database Connectivity (JDBC), LDAP datastores, or more sophisticated JDBC or LDAP queries, including flat files or SOAPconnected databases. You can write datastores to perform configuration assistance or validation actions, such as testing a connection to a database. Actions can also include generation of parameters that might need manual setting in a configuration file.

For more information, see the Javadoc for the CustomDataSourceDriver interface, the SamplePropertiesDataStore.java file for a sample implementation, and the SDK Developer's Guide on page 1021 for build and deployment information.



The Javadoc for PingFederate and the sample implementation are in the <pf install>/ pingfederate/sdk directory.

Steps

1. After the data store driver (JAR) file is written and installed, select it in the **Data Store** window when creating a new instance of your data store.

2. On the **Configure Data Store Instance** tab, configure your data store instance.



Depending on the data store implementation, configuration requirements vary.

Example:

After building and deploying the sample from the <pf install>/pingfederate/sdk/pluginsrc/custom-data-store-example directory, you can create an instance of the Sample SDK Properties Data Store and configure the rest, as shown in the following images.



Note:

When editing an existing instance, you can modify the name of the data store instance and toggle the option for if PingFederate should mask attribute values returned from this data store instance in PingFederate logs.

3. On the Actions window, follow the on-screen instructions provided by the developer to complete any required tasks.



Note:

Depending on the datastore implementation, configuration requirements vary. If no action is required, this window is not shown.

4. Click **Save** to keep your configuration.

Defining a datastore for persistent authentication sessions

When enabling PingFederate authentication sessions, you can select the persistent option so that PingFederate can leverage previous sessions as users request protected resources after restarting their browsers.

About this task

This optional persistent configuration requires external storage of session-state data, as opposed to in-memory alone. By default, PingFederate uses its internal HSQLDB database to maintain persistent authentications. You can configure PingFederate to maintain persistent authentication sessions externally on a database server or a PingDirectory server.



CAUTION:

Use the built-in HSQLDB only for trial or training environments. For testing and production environments, always use a secured external storage solution for proper functioning in a clustered environment.

Testing involving HSQLDB is not a valid test. In both testing and production, it might cause various problems due to its limitations and HSQLDB involved cases are not supported by Pingldentity.

Steps

- 1. Create the required data structure on the external storage medium.
- 2. Modify two PingFederate configuration XML files.

Configuring an external database for authentication sessions

Set up various tables so that PingFederate can store authentication sessions on corresponding database servers.

About this task

Specific tables are required in order for PingFederate to store authentication sessions on your database server. Table-setup scripts are provided for supported database servers.

Steps

- 1. Run the table-setup scripts, provided in the <pf_install>/pingfederate/server/default/conf/authentication-session/sql-scripts directory, for your database server.
- 2. If you have not already done so, go to **System # Data & Credential Stores**. In the **Data Stores** window, create a Java Database Connection (JDBC) datastore for your database server.
- 3. Copy the system ID of the applicable JDBC datastore from the **Data Stores** window.
- 4. Edit the

org.sourceid.saml20.service.session.data.impl.SessionStorageManagerJdbcImpl.xml file, located in the $<\!pf_install>\!$ /pingfederate/server/default/data/config-store directory.

Replace the <c:item name="PingFederateDSJNDIName"/> element value with the system ID of your data store connection and save the file.

Example:

For example, if the system ID is JDBC-123456789ABCDEF123456789ABCDEF123456A0A6, update the

org.sourceid.saml20.service.session.data.impl.SessionStorageManagerJdbcImpl.xml file as follows.

5. Edit the <pf_install>/pingfederate/server/default/conf/META-INF/hivemodule.xml file.

Example: If the system ID is SessionStorageManager, locate the SessionStorageManager service point, set the value of the class attribute to org.sourceid.saml20.service.session.data.impl.SessionStorageManagerJdbcImpl (the default value), then save the file.

</service-point>



Note:

For a clustered PingFederate environment, you must edit the hivemodule.xml file on each node manually as cluster replication cannot replicate this change to other nodes.

6. Start or restart the PingFederate service.



For a clustered PingFederate environment, replicate this new configuration to other engine nodes on System # Server # Cluster Management. Start or restart the PingFederate service on each engine node to activate the change.

Result

PingFederate removes expired authentication sessions from the database once a day. To fine-tune the frequency and the number of expired authentication sessions to remove, see *Managing authentication* sessions stored in the database on page 821.

Configuring PingDirectory for authentication sessions

Use specific schema objects in order to enable PingFederate to store authentication sessions on your directory server. For PingDirectory, LDIF scripts are provided for this purpose.

Steps

- 1. Update the LDAP schema.
 - a. Sign on to the PingDirectory administrative console.
 - b. Go to LDAP Schema # Schema Utilities.
 - c. Click Import Schema Element.
 - d. Copy the schema changes from the authentication-session-attributes-ldappingdirectory.ldif file and paste them into the text area.

The file is located in the <pf install>/pingfederate/server/default/conf/ authentication-session/ldif-scripts directory.

Replace the placeholder values with relevant information from your directory server.

- e. Click **Import**.
- 2. Create the following indexes.

Attribute name	Index type
pf-authn-session-group-hashed-session-id	equality
pf-authn-session-group-user-ids	equality
pf-authn-session-group-expiry-time	ordering
pf-authn-session-group-last-activity-time	ordering

Create these indexes with Ping Directory's dsconfig utility. The dsconfig utility is interactive. You can also provide inputs as command arguments. The following examples create the indexes.

```
$ bin/dsconfig create-local-db-index \
  --backend-name userRoot \
  --index-name pf-authn-session-group-hashed-session-id \
```

```
--set index-type:equality
$ bin/dsconfig create-local-db-index \
 --backend-name userRoot \
 --index-name pf-authn-session-group-user-ids \
 --set index-type:equality
$ bin/dsconfig create-local-db-index \
 --backend-name userRoot \
 --index-name pf-authn-session-group-expiry-time \
 --set index-type:ordering
$ bin/dsconfig create-local-db-index \
 --backend-name userRoot \
 --index-name pf-authn-session-group-last-activity-time \
 --set index-type:ordering
```

After adding the indexes, use the rebuild-index utility to build the indexes. The following example builds the required indexes.

```
$ bin/rebuild-index \
  --baseDN "dc=example, dc=com" \
  --index pf-authn-session-group-hashed-session-id \
  --index pf-authn-session-group-user-ids \
  --index pf-authn-session-group-expiry-time \
  --index pf-authn-session-group-last-activity-time
```

For more information, see Working with Indexes in the PingDirectory Administration Guide .

- 3. If you have not already done so, create an LDAP data store for your directory server on **System** # Data & Credential Stores # Data Stores.
- 4. Copy the system ID of the applicable LDAP data store from the **Data Stores** window.
- 5. Edit the /pingfederate/server/default/data/config-store/ org.sourceid.saml20.service.session.data.impl.SessionStorageManagerLdapImpl.xml file.



Note:

For a clustered PingFederate environment, edit this file on the administrative console node first, and then replicate to other engine nodes using System # Server # Cluster Management.

a. Replace the <c:item name="PingFederateDSJNDIName"/> element value with the system ID of your data store connection.

Example:

For example, if the system ID is LDAP-123456789ABCDEF123456789ABCDEF123456A0AC, update the configuration file as follows.

```
<!-- Data store id -->
name="PingFederateDSJNDIName">LDAP-123456789ABCDEF123456789ABCDEF123456A0AC</
```

b. Enter a value for the <c:item name="SearchBase"/> element.



This is the distinguished name (DN) that points to the client location. For more information, see the inline comment and the LDIF scripts in the <pf install>/pingfederate/server/ default/conf/authentication-session/ldif-scripts directory.

- c. Update the attribute names only if you have changed attribute names in the LDIF scripts located in the f install>/pingfederate/server/default/conf/authentication-session/ ldif-scripts directory.
- d. Save the file.
- 6. Edit the < pf install > / pingfederate/server/default/conf/META-INF/hivemodule.xml
 - a. Locate the SessionStorageManager service point.

```
<!-- Service for storing Authentication Sessions. -->
<service-point id="SessionStorageManager"</pre>
 interface="org.sourceid.saml20.service.session.data.SessionStorageManager">
    <invoke-factory>
        <!--
        Supported classes are
org.sourceid.saml20.service.session.data.impl.SessionStorageManagerJdbcImpl :
Use this service-point for a Jdbc implementation.
org.sourceid.saml20.service.session.data.impl.SessionStorageManagerLdapImpl :
Use this service-point for an LDAP implementation.
        <construct
class="org.sourceid.saml20.service.session.data.impl.SessionStorageManagerJdbcIm
    </invoke-factory>
</service-point>
```

b. Set the value of the class attribute to

org.sourceid.saml20.service.session.data.impl.SessionStorageManagerLdapImpl.

c. Save the file.



For a clustered PingFederate environment, you must edit the hivemodule.xml file on each node manually as cluster replication cannot replicate this change to other nodes.



For a clustered PingFederate environment, replicate this new configuration to other engine nodes on System # Server # Cluster Management. Start or restart the PingFederate service on each engine node to activate the change.



Note:

When storing persistent authentication sessions on a PingDirectory server, you must also configure a cleanup plugin in PingDirectory to remove expired authentication sessions from your directory server. For more information, see *Managing authentication sessions stored in PingDirectory* on page 824.

OAuth grant datastores

Learn about persistent grant data stores and persistent authorizations.

PingFederate uses a built-in HSQLDB as its persistent grant datastore after the initial setup.



CAUTION:

Use the built-in HSQLDB only for trial or training environments. For testing and production environments, always use a secured external storage solution for proper functioning in a clustered environment.

Testing involving HSQLDB is not a valid test. In both testing and production, it may cause various problems due to its limitations and HSQLDB involved cases are not supported by Pingldentity.

Persistent authorizations include those obtained by OAuth clients in the following ways:

Grants obtained or updated using the authorization code, resource owner credentials, or device authorization grant type, in conjunction with the refresh token grant type



Note:

If the use cases involve mapping attributes from authentication sources, such as IdP adapter instances or IdP connections, or password credential validator (PCV) instances to the access tokens, directly or through persistent grant-extended attributes, storing these attributes from authentication sources and their values along with the persistent grants maintains them for reuse when clients subsequently present refresh tokens for new access tokens.

Grants obtained or updated by using the implicit grant type, for which PingFederate is configured to reuse existing persistent grants



Note:

If the use cases involve mapping attributes from authentication sources or PCV instances to the access tokens, runtime procedures obtain attribute values for each token request, but persistent grants do not store with attributes or their values.

Persistent grants, and any associated attributes and their values, remain valid until the grants expire or until PingFederate explicitly revokes them or cleans them up.



Note:

Attribute values are always stored encrypted when a directory is used. If a database server is used (including the built-in HSQLDB database), attribute values are also stored encrypted by default.

Configuring external databases for grant storage

Specific tables are required in order for PingFederate to store grants, the associated attributes, and their values (if any), on your database server. Table-setup scripts are provided for supported database servers.

Steps

- 1. Run the table-setup scripts for your database server provided in the <pf install>/ pingfederate/server/default/conf/access-grant/sgl-scripts directory.
- 2. If you have not already done so, create a JDBC data store for your database server on System # Data & Credential Stores # Data Stores.
- 3. Copy the system ID of the applicable Java Database Connection (JDBC) data store from the Data Stores window.
- 4. Edit the cpf install>/pingfederate/server/default/data/config-store/ org.sourceid.oauth20.token.AccessGrantManagerJdbcImpl.xml file.

Replace the <c:item name="PingFederateDSJNDIName"/> element value with the system ID of your data store connection and save the file.

Example:

If the system ID is JDBC-123456789ABCDEF123456789ABCDEF123456A0A6, update the org.sourceid.oauth20.token.AccessGrantManagerJdbcImpl.xml file as follows.

```
<?xml version="1.0" encoding="UTF-8"?>
<c:config xmlns:c="http://www.sourceid.org/2004/05/config">
    <c:item
name="PingFederateDSJNDIName">JDBC-123456789ABCDEF123456789ABCDEF123456A0A6</
c:item>
</c:config>
```

- 5. Edit the <pf install>/pingfederate/server/default/conf/META-INF/hivemodule.xml
 - a. Locate the AccessGrantManager service point.

```
<!-- Service for storage of access grants -->
<service-point id="AccessGrantManager"</pre>
interface="com.pingidentity.sdk.accessgrant.AccessGrantManager">
    <create-instance</pre>
class="org.sourceid.oauth20.token.AccessGrantManagerJdbcImpl"/>
</service-point>
```

- b. Set the value of the class attribute to org.sourceid.oauth20.token.AccessGrantManagerJdbcImpl, the default value.
- c. Save the file.



Note:

For a clustered PingFederate environment, you must edit the hivemodule.xml file on each node manually as cluster replication cannot replicate this change to other nodes.



For a clustered PingFederate environment, replicate this new configuration to other engine nodes on System # Server # Cluster Management. Start or restart the PingFederate service on each engine node to activate the change.

Result

PingFederate provides two cleanup tasks for persistent grants. One task manages expired grants, while another task caps the number of grants based on a combination of user, client, grant type, and authentication context. For more information, see OAuth persistent grants cleanup on page 826.

Configuring directories for grant storage

PingFederate requires specific schema objects in order to store grants, the associated attributes, and their values (if any) on your directory server. LDIF scripts are provided for supported directory servers.

Steps

- 1. Review the LDIF scripts for your directory server provided in the <pf install>/pingfederate/ server/default/conf/access-grant/ldif-scripts directory.
- 2. Replace placeholder values with relevant information from your directory server.
- 3. Run the LDIF scripts to update your LDAP schema.



Note:

For Active Directory, run the script to create the attributes, then run the script to create the object class.

- 4. If you have not already done so, create an LDAP datastore for your directory server on System # Data & Credential Stores # Data Stores.
- 5. Copy the system ID of the applicable LDAP datastore from System # Data & Credential Stores # Data Stores.
- 6. Edit the configuration file relevant to your directory server.



Note:

This configuration file is located in the <pf install>/pingfederate/server/default/data/ config-store directory, as described in the following table.

Directory server	Configuration file	
PingDirectory	org.sourceid.oauth20.token.AccessGrantManagerLDAPP	ingDirecto
Microsoft Active Directory	org.sourceid.oauth20.token.AccessGrantManagerLDAPA	DImpl.xml

Copyright ©2024

Directory server	Configuration file	
Oracle Directory Server Enterprise Edition or Oracle Unified Directory	org.sourceid.oauth20.token.AccessGrantManagerLDAPO	racleImpl

a. Replace the <c:item name="PingFederateDSJNDIName"/> element value with the system ID of your datastore connection.

Example:

If the system ID is LDAP-123456789ABCDEF123456789ABCDEF123456A0A6, update the configuration file as follows.

```
<!-- Data store id -->
<c:item
name="PingFederateDSJNDIName">LDAP-123456789ABCDEF123456789ABCDEF123456A0A6</
c:item>
```

b. Enter a value for the <c:item name="SearchBase"/> element.



This is the distinguished name (DN) that points to the access grants location. For more information, see the inline comment and the LDIF scripts in the <pf install>/ pingfederate/server/default/conf/access-grant/ldif-scripts directory.

- c. Update the attribute names only if you have changed attribute names in the LDIF scripts located in the <pf install>/pingfederate/server/default/conf/access-grant/ldifscripts directory.
- d. Save the file.
- 7. Edit the <pf install>/pingfederate/server/default/conf/META-INF/hivemodule.xml
 - a. Locate the AccessGrantManager service point.

```
<!-- Service for storage of access grants -->
<service-point id="AccessGrantManager"</pre>
interface="com.pingidentity.sdk.accessgrant.AccessGrantManager">
     <create-instance</pre>
class="org.sourceid.oauth20.token.AccessGrantManagerJdbcImpl"/>
</service-point>
```

b. Update the class attribute value to one of the following values.

Directory server	Class value	
PingDirectory	org.sourceid.oauth20.token.AccessGrantManagerLDA	PPingDire
Microsoft Active Directory	org.sourceid.oauth20.token.AccessGrantManagerLDA	PADImpl

Directory server	Class value	
Oracle Directory Server Enterprise Edition or Oracle Unified Directory	org.sourceid.oauth20.token.AccessGrantManagerLDA	.POracleImj

c. Save the file.



Note:

For a clustered PingFederate environment, you must edit the hivemodule.xml file on each node manually as cluster replication cannot replicate this change to other nodes.

8. Start or restart PingFederate.



For a clustered PingFederate environment, replicate this new configuration to other engine nodes on System # Server # Cluster Management. Start or restart the PingFederate service on each engine node to activate the change.

9. In the directory, create indexes for the following OAuth grant attributes. If you are using PingDirectory, see Indexing grant attributes in PingDirectory on page 893 for more information.

Attribute name	Index type
accessGrantGuid	equality
accessGrantUniqueUserIdentifier	equality
accessGrantHashedRefreshTokenValue	equality
accessGrantClientId	equality
accessGrantExpires	ordering

Result

PingFederate provides two cleanup tasks for persistent grants. One task manages expired grants, while another task caps the number of grants based on a combination of user, client, grant type, and authentication context. For more information, see OAuth persistent grants cleanup on page 826.

Indexing grant attributes in PingDirectory

If you use PingDirectory, or another directory, to store OAuth persistent grants for PingFederate, you must index the grant attributes.

About this task

Index these OAuth grant attributes using the procedure below.

Attribute name	Index type
accessGrantGuid	equality
accessGrantUniqueUserIdentifier	equality
accessGrantHashedRefreshTokenValue	equality
accessGrantClientId	equality
accessGrantExpires	ordering

1. Create the indexes using PingDirectory's dsconfig utility.

The dsconfig utility is interactive, letting you enter command arguments. The following examples create the required indexes.

```
$ bin/dsconfig create-local-db-index \
  --backend-name userRoot \
  --index-name accessGrantGuid \
  --set index-type:equality
$ bin/dsconfig create-local-db-index \
  --backend-name userRoot \
  --index-name accessGrantUniqueUserIdentifier \
  --set index-type:equality
$ bin/dsconfig create-local-db-index \
 --backend-name userRoot \
  --index-name accessGrantHashedRefreshTokenValue \
  --set index-type:equality
$ bin/dsconfig create-local-db-index \
  --backend-name userRoot \
  --index-name accessGrantClientId \
  --set index-type:equality
$ bin/dsconfig create-local-db-index \
 --backend-name userRoot \
 --index-name accessGrantExpires \
  --set index-type:ordering
```

2. After adding the indexes, build the indexes using the rebuild-index utility.

The following example builds the required indexes.

```
$ bin/rebuild-index \
 --baseDN "dc=example,dc=com" \
  --index accessGrantGuid \
 --index accessGrantUniqueUserIdentifier \
 --index accessGrantHashedRefreshTokenValue \
  --index accessGrantClientId \
  --index accessGrantExpires
```



Note:

You can configure a PingDirectory plugin to handle the cleanup of expired persistent grants and the associated attributes. The plugin allows fine-grained control over various aspects of the cleanup task, which can smooth out the performance impact. For more information, see *Managing expired persistent* grants in PingDirectory on page 828.

Using custom solutions for grant storage

Use the PingFederate SDK to implement a custom solution for OAuth grant storage.

Steps

1. Implement the AccessGrantManager interface.



For more information, see the Javadoc for the AccessGrantManager interface, the SampleAccessGrant.java file for a sample implementation, and the SDK Developer's Guide on page 1021 for build and deployment information.



Tip:

The Javadoc for PingFederate and the sample implementation are in the <pf install>/ pingfederate/sdk directory.

- 2. Edit the <pf install>/pingfederate/server/default/conf/META-INF/hivemodule.xml
 - a. Locate the AccessGrantManager service point.

```
<!-- Service for storage of access grants -->
<service-point id="AccessGrantManager"</pre>
interface="com.pingidentity.sdk.accessgrant.AccessGrantManager">
     <create-instance</pre>
class="org.sourceid.oauth20.token.AccessGrantManagerJdbcImpl"/>
</service-point>
```

- b. Update the class attribute value to the name of your class.
- c. Save the file.



Note:

For a clustered PingFederate environment, you must edit the hivemodule.xml file on each node manually as cluster replication cannot replicate this change to other nodes.

- 3. Deploy the required program files of your custom implementation to all PingFederate servers.
- 4. Start or restart PingFederate.



Note:

For a clustered PingFederate environment, replicate this new configuration to other engine nodes on System # Server # Cluster Management. Start or restart the PingFederate service on each engine node to activate the change.

OAuth client datastores

Change the default storage method of XML files in PingFederate in order to make it easier to register clients or manage their records through the OAuth Client Management Service.

PingFederate stores client records in XML files by default. On-disk storage allows you to manage clients using the administrative console and the administrative API. Client records are part of the configuration archive.

You can configure PingFederate to store client records externally, which provides the flexibility to manage client records through the OAuth Client Management Service, or enable dynamic client registration for your partner-developers. In this scenario, client records are not part of the configuration archive. They are stored on a database server, a directory server, or some other storage medium through the use of the PingFederate SDK.

Configuring external databases for client storage

Specific tables are required in order for PingFederate to store OAuth client records on your database server. Table-setup scripts are provided for supported database servers.

About this task



CAUTION:

PingFederate does not migrate client records from one storage medium to another. You must recreate your clients after updating the client storage configuration. If you need only a few clients, you can recreate them using the administrative console.

If you need a large number of clients, use the administrative API to retrieve your client records before updating the client storage. Update the client storage configuration and recreate your clients using the administrative API based on the retrieved records. For more information, see .



CAUTION:

A pre-installed, default HSQLDB database is selected for initial setup and testing. However, we strongly recommend that you use the built-in HSQLDB only for trial or training environments. For testing and production environments, always use a secured external storage solution for proper functioning in a clustered environment.

Testing involving HSQLDB is not a valid test. In both testing and production, it may cause various problems due to its limitations and HSQLDB involved cases are not supported by Pingldentity.

Steps

- 1. Edit the <pf install>/pingfederate/server/default/conf/META-INF/hivemodule.xml
 - a. Locate the ClientManager service point.

```
<!-- Service for storing OAuth client configuration. -->
<service-point id="ClientManager"</pre>
 interface="org.sourceid.oauth20.domain.ClientManager">
    <invoke-factory>
       <!--
        Supported classes are
        org.sourceid.oauth20.domain.ClientManagerXmlFileImpl ...
        org.sourceid.oauth20.domain.ClientManagerJdbcImpl
        org.sourceid.oauth20.domain.ClientManagerLdapImpl
        org.sourceid.oauth20.domain.ClientManagerGenericImpl ...
        -->
        <construct
 class="org.sourceid.oauth20.domain.ClientManagerXmlFileImpl"/>
    </invoke-factory>
```

```
</service-point>
```

- b. Update the class attribute value to org.sourceid.oauth20.domain.ClientManagerJdbcImpl.
- c. Save the file.



Important:

You must set up an external database because you cannot share the bundled HSQLDB database across multiple PingFederate engine nodes. For production standalone deployments, we recommend you store the client records in an external secured database.

- 2. Run the table-setup scripts for your database server provided in the <pf install>/ pingfederate/server/default/conf/oauth-client-management/sql-scripts directory.
- 3. If you have not already done so, create a Java Database Connectivity (JDBC) datastore for your database server. Go to = **System # Data & Credential Stores**.
- 4. In the **Data Stores** window, copy the system ID of the applicable JDBC datastore.
- 5. Edit the <pf install>/pingfederate/server/default/data/config-store/ org.sourceid.oauth20.domain.ClientManagerJdbcImpl.xml file.

Replace the <c:item name="PingFederateDSJNDIName"/> element value with the system ID of your datastore connection and save the file.

Example:

If the system ID is JDBC-123456789ABCDEF123456789ABCDEF123456A0AC, update the org.sourceid.oauth20.domain.ClientManagerJdbcImpl.xml file as follows.

```
<?xml version="1.0" encoding="UTF-8"?>
<c:config xmlns:c="http://www.sourceid.org/2004/05/config">
    <c:item
name="PingFederateDSJNDIName">JDBC-123456789ABCDEF123456789ABCDEF123456A0AC</
c:item>
</c:config>
```

6. Start or restart PingFederate.



For a clustered PingFederate environment, replicate this new configuration to other engine nodes on System # Server # Cluster Management. Start or restart the PingFederate service on each engine node to activate the change.

Configuring directories for client storage

Specific schema objects are required in order for PingFederate to store OAuth client records on your directory server. LDIF scripts are provided for supported directory servers.

About this task



CAUTION:

PingFederate does not migrate client records from one storage medium to another. You must recreate your clients after updating the client storage configuration. If you need only a few clients, you can recreate them using the administrative console.

If you need a large number of clients, use the administrative API to retrieve your client records before updating the client storage. Update the client storage configuration and recreate your clients using the administrative API based on the retrieved records. For more information, see .

Steps

- 1. Review the LDIF scripts for your directory server provided in the <pf install>/pingfederate/ server/default/conf/oauth-client-management/ldif-scripts directory.
- 2. Replace placeholder values with relevant information from your directory server.
- 3. Run the LDIF scripts to update your LDAP schema.



Note:

For Active Directory, run the script to create the attributes, then run the script to create the object

- 4. If you have not already done so, create an LDAP datastore for your directory server on System # Data & Credential Stores # Data Stores.
- 5. Copy the system ID of the applicable LDAP datastore from System # Data & Credential Stores # Data Stores.
- 6. Edit the <pf install>/pingfederate/server/default/data/config-store/ org.sourceid.oauth20.domain.ClientManagerLdapImpl.xml file.
 - a. Replace the <c:item name="PingFederateDSJNDIName"/> element value with the system ID of your datastore connection.

Example:

If the system ID is LDAP-123456789ABCDEF123456789ABCDEF123456A0AC, update the configuration file as follows.

```
<!-- Data store id -->
<c:item
name="PingFederateDSJNDIName">LDAP-123456789ABCDEF123456789ABCDEF123456A0AC</
c:item>
. . .
```

b. Enter a value for the <c:item name="SearchBase"/> element.



This is the distinguished name (DN) that points to the client location. For more information, see the inline comment and the LDIF scripts in the <pf install>/pingfederate/server/ default/conf/oauth-client-management/ldif-scripts directory.

- c. Update the attribute names only if you have changed attribute names in the LDIF scripts located in the <pf install>/pingfederate/server/default/conf/oauth-clientmanagement/ldif-scripts directory.
- d. Save the file.
- 7. Edit the <pf install>/pingfederate/server/default/conf/META-INF/hivemodule.xml
 - a. Locate the ClientManager service point.

```
<!-- Service for storing OAuth client configuration. -->
<service-point id="ClientManager"</pre>
interface="org.sourceid.oauth20.domain.ClientManager">
```

```
<invoke-factory>
        <!--
        Supported classes are
        org.sourceid.oauth20.domain.ClientManagerXmlFileImpl ...
        org.sourceid.oauth20.domain.ClientManagerJdbcImpl
        org.sourceid.oauth20.domain.ClientManagerLdapImpl
        org.sourceid.oauth20.domain.ClientManagerGenericImpl ...
        -->
       <construct
class="org.sourceid.oauth20.domain.ClientManagerXmlFileImpl"/>
   </invoke-factory>
</service-point>
```

- b. Update the class attribute value to
 - org.sourceid.oauth20.domain.ClientManagerLdapImpl.
- c. Save the file.



Note:

For a clustered PingFederate environment, you must edit the hivemodule.xml file on each node manually as cluster replication cannot replicate this change to other nodes.

8. Start or restart PingFederate.



For a clustered PingFederate environment, replicate this new configuration to other engine nodes on System # Server # Cluster Management. Start or restart the PingFederate service on each engine node to activate the change.

9. In the directory, create indexes for the following OAuth client attributes. If you are using PingDirectory, see Indexing client attributes in PingDirectory on page 899 for more information.

Attribute name	Index type
pf-oauth-client-id	equality
pf-oauth-client-id	ordering
pf-oauth-client-id	substring
pf-oauth-client-name	equality
pf-oauth-client-name	ordering
pf-oauth-client-name	substring
pf-oauth-client-last-modified	ordering

Indexing client attributes in PingDirectory

If you use PingDirectory, or another directory, to store OAuth client records for PingFederate, you must index the client attributes.

About this task

Index these OAuth client attributes using the procedure below.

Attribute name	Index type
pf-oauth-client-id	equality
pf-oauth-client-id	ordering

Attribute name	Index type
pf-oauth-client-id	substring
pf-oauth-client-name	equality
pf-oauth-client-name	ordering
pf-oauth-client-name	substring
pf-oauth-client-last-modified	ordering

Steps

1. Create the indexes using the PingDirectory dsconfig utility.

The dsconfig utility is interactive, letting you enter command arguments. The following example creates the three indexes for the pf-oauth-client-id attribute.

```
$ bin/dsconfig create-local-db-index \
 --backend-name userRoot \
  --index-name pf-oauth-client-id \
  --set index-type:equality \
  --set index-type:ordering \
  --set index-type:substring
```

2. After creating the indexes, build them using the rebuild-index utility.

The following example builds the required indexes.

```
$ bin/rebuild-index \
 --baseDN "dc=example,dc=com" \
 --index pf-oauth-client-id \
 --index pf-oauth-client-name \
 --index pf-oauth-client-last-modified
```

Using custom solutions for client storage

Use the PingFederate SDK to implement a custom solution for OAuth client storage.

About this task



CAUTION:

PingFederate does not migrate client records from one storage medium to another. You must recreate your clients after updating the client storage configuration. If you need only a few clients, you can recreate them using the administrative console.

If you need a large number of clients, use the administrative API to retrieve your client records before updating the client storage. Update the client storage configuration and recreate your clients using the administrative API based on the retrieved records. For more information, see .

Steps

1. Implement the ClientStorageManagerV2 interface.

This interface includes a search () method, allowing developers to provide efficient implementations of the pagination and search functions exposed in the administrative console.

For more information, see the Javadoc for the ClientStorageManagerV2 interface, the SampleClientStorage.java file for a sample implementation, and the SDK Developer's Guide on page 1021 for build and deployment information.



Tip:

The Javadoc for PingFederate and the sample implementation are in the <pf install>/ pingfederate/sdk directory.

- $\textbf{2. Edit the} < \!\!\! \textit{pf install} \!\!\! > \!\!\!\! \texttt{pingfederate/server/default/conf/META-INF/hivemodule.xml}$
 - a. Locate the ClientStorageManager service point.

```
<!-- Service for storing OAuth client configuration. -->
<service-point id="ClientManager"</pre>
interface="org.sourceid.oauth20.domain.ClientManager">
   <invoke-factory>
        <!--
        Supported classes are
        org.sourceid.oauth20.domain.ClientManagerXmlFileImpl ...
        org.sourceid.oauth20.domain.ClientManagerJdbcImpl
        org.sourceid.oauth20.domain.ClientManagerLdapImpl
        org.sourceid.oauth20.domain.ClientManagerGenericImpl ...
        <construct
class="org.sourceid.oauth20.domain.ClientManagerXmlFileImpl"/>
    </invoke-factory>
</service-point>
```

- b. Update the class attribute value with the name of the class implementing the ClientStorageManagerV2 interface.
- c. Save the file.



Note:

For a clustered PingFederate environment, you must edit the hivemodule.xml file on each node manually as cluster replication cannot replicate this change to other nodes.

3. Start or restart PingFederate.



Note:

For a clustered PingFederate environment, replicate this new configuration to other engine nodes on System # Server # Cluster Management. Start or restart the PingFederate service on each engine node to activate the change.

Account-linking datastores

Configure where you want to store account links, either internally or externally.

When a service provider (SP) is configured to use account linking for an identity provider (IdP) connection, by default PingFederate uses the built-in Hyper SQL Database (HSQLDB) as the account-link repository.

You can also configure PingFederate to store account links on an external database server or directory server. For specific instructions on how to configure these options, see the following topics:

- Configuring external databases for account-link storage on page 902
- Configuring directories for account-link storage on page 902



CAUTION:

Use the built-in HSQLDB only for trial or training environments. For testing and production environments, always use a secured external storage solution for proper functioning in a clustered environment.

Testing involving HSQLDB is not a valid test. In both testing and production, it might cause various problems due to its limitations and HSQLDB involved cases are not supported by Pingldentity.

Configuring external databases for account-link storage

A specific table is required in order for PingFederate to store account links on your database server. Tablesetup scripts are provided for supported database servers.

Steps

- 1. Create a database for account linking using one of the table-setup scripts located in the <pf install>/pingfederate/server/default/conf/account-linking/sql-scripts directory.
- 2. Go to System # Data & Credential Stores # Data Stores and create a new datastore to connect PingFederate to the database. For more information, see Configuring a JDBC connection on page 868.
- 3. On the Data Stores window, copy the system ID of the new account-linking datastore.
- 4. In the org.sourceid.saml20.service.impl.AccountLinkingServiceDBImpl.xml file, located in the <pf install>/pingfederate/server/default/data/config-store directory, replace the <c:item name="PingFederateDSJNDIName"/> element value with the system ID of your datastore connection and save the file.

Example:

For example, if the system ID is JDBC-123456789ABCDEF123456789ABCDEF123456A0AC, update the org.sourceid.saml20.service.impl.AccountLinkingServiceDBImpl.xml file as follows.

```
<?xml version="1.0" encoding="UTF-8"?>
<c:config xmlns:c="http://www.sourceid.org/2004/05/config">
    <c:item
name="PingFederateDSJNDIName">JDBC-123456789ABCDEF123456789ABCDEF123456A0AC</
c:item>
</c:config>
```

- 5. Start or restart PingFederate.
- 6. If you are running PingFederate in a cluster, go to System # Server # Cluster Management and replicate this change to other runtime servers.

Configuring directories for account-link storage

You can create and configure a directory server to store account linking data.

Before you begin



Note:

User accounts for linking must exist in the directory prior to establishing the account link. The Account Linking Service does not add users to the directory server, it only updates AccountLinkDataAttribute for a given user.

Steps

- Go to System # Data & Credential Stores # Data Stores and create a new datastore to connect PingFederate to the directory. For more information, see Configuring an LDAP connection on page 871
- 2. Copy the system ID of the new account-linking datastore.
- 3. Edit the hivemodule.xml file.

The file is located in the $<\!pf_install>\!/pingfederate/server/default/conf/META-INF directory.$

Locate the service-point for the Account Linking Service, as shown in the following example.

Update the class value to

org.sourceid.saml20.service.impl.AccountLinkingServiceLDAPImpl, as shown in the following example.

Locate the Service-Point ID for AccountLinkingService and change the value of the create-instance class to org.sourceid.saml20.service.impl.AccountLinkingServiceLDAPImpl.

4. Edit the <pf_install>/
org.sourceid.saml20.service.impl.AccountLinkingServiceLDAPImpl.xml file.

The following example shows the default content of the file.

```
<?xml version="1.0" encoding="UTF-8"?>
<c:config xmlns:c="http://www.sourceid.org/2004/05/config">
```

```
<!-- Data store id -->
    <c:item name="PingFederateDSJNDIName"></c:item>
    <!-- LDAP search base -->
    <c:item name="UserSearchBase"></c:item>
    <!-- LDAP username attribute. ex: sAMAccountName -->
    <c:item name="UsernameAttribute"></c:item>
    <!-- Attribute on user object to place Account Linking data -->
    <c:item name="AccountLinkDataAttribute"></c:item>
</c:config>
```

Insert the applicable values between the XML tags as shown in the following table.

Item name Element value PingFederateDSJNDINahe system ID of new account-linking datastore. The location in the directory server from which the search begins. UserSearchBase UsernameAttribute The attribute that represents the user identifier. AccountLinkDataAttrIbe attribute to store account linking data. Note: The AccountLinkDataAttribute can be any multivalued string attribute on a user object class. We recommend that you extend the LDAP schema with a custom attribute for use here. For more information on extending the Active Directory schema, see Extending the Schema from Microsoft.

- 5. Start or restart PingFederate.
- 6. If you are running PingFederate in a cluster, go to System # Server # Cluster Management and replicate this change to the other runtime servers.



Note:

You must also manually apply the changes made in the hivemodule.xml file and then start or restart PingFederate on each runtime server.

7. In the directory, create equality indexes on the LDAP attribute types you specified for the configuration properties UsernameAttribute and AccountLinkDataAttribute.

Example: For example, you would need to create equality indexes on sAMAccountName and AccountLink if you had specified the following in step 4:

```
<!-- LDAP username attribute. ex: sAMAccountName -->
<c:item name="UsernameAttribute">sAMAccountName</c:item>
<!-- Attribute on user object to place Account Linking data -->
 <c:item name="AccountLinkDataAttribute">AccountLink</c:item>
```

PingFederate provides an authentication mechanism using plugin password credential validators (PCVs). This feature provides centralized credential validation for various PingFederate components and configurations.

To manage Password Credential Validators, go to System # Data & Credential Stores # Password Credential Validators.

For each instance of the HTML Form Adapter, the HTTP Basic Adapter, and the Username Token Processor, you can select the same PCV instance, a unique PCV instance, or multiple PCV instances. When you select multiple PCV instances for a given adapter or token processor instance, if the first PCV instance fails to authenticate a user, the PCV returns control to the adapter or the token processor. The adapter or the token processor then tries the next PCV instance. The cycle stops until a PCV instance succeeds or the last PCV instance also fails.

For OAuth clients using the Resource Owner Password Credentials grant type, you configure a grantmapping configuration to fulfill the persistent grant contract using the attribute values from the applicable PCV instances.



Note:

You can only create one grant-mapping configuration per applicable PCV instance.

If you want to manage OAuth client records using the OAuth Client Management Service or persistent grants using the OAuth Access Grant Management Service, you must select a PCV instance when configuring authorization server settings. When accessing these services, you must include in the requests valid credentials via HTTP Basic authentication scheme.

PingFederate is distributed with the following plugin PCVs.

LDAP Username Password Credential Validator

Validates credentials based on an LDAP look-up in an organization's user-datastore.

PingID PCV (with integrated RADIUS server)

Validates credentials from a VPN RADIUS client based on an LDAP look-up in an organization's user-datastore. For more information, see Integrate PingID with your VPN.

PingOne for Enterprise Directory Password Credential Validator

Validates credentials stored in PingOne for Enterprise Directory.

RADIUS Username Password Credential Validator

Validates credentials based on the RADIUS protocol on an organization's RADIUS server.

Simple Username Password Credential Validator

Validates credentials maintained by PingFederate.



By default, PingFederate automatically checks multi-connection errors whenever you access this window. This verifies that configured connections are not adversely affected by changes made here.

If you experience noticeable delays in accessing this window, you can disable automatic connection validation. Go to System # Server # General Settings.

Choosing a Password Credential Validator

Choose the type of Password Credential Validator (PCV) you will use in PingFederate. You must also specify the PCV's name, ID, and whether it uses a parent instance.

About this task

Available PCV types are determined by plug-in .jar files loaded in the $<\!pf_install>/pingfederate/server/default/deploy directory.$ Several validator plugins are bundled with PingFederate. You can add other plugins from the Ping Identity Downloads website.

Steps

- 1. On the **Type** tab, enter a name and an ID for the instance.
- 2. Select the type of the PCV from the **Type** list.
- 3. Optional: Select a Parent Instance from the list.

Use this option when creating an instance that is similar to an existing one. The child instance inherits the configuration of its parent. You can also override one or more settings during the setup. Select the **Override** ... check box and make the adjustments as needed in one or more subsequent windows.

Password Credential Validator instance configurations

The instance configuration of a Password Credential Validator (PCV) varies depending on the credential validators deployed on your server.

For PCVs bundled with PingFederate, see the following topics:

- Configuring the LDAP Username Password Credential Validator on page 906
- Configuring the PingOne for Enterprise Directory Password Credential Validator on page 911
- Configuring the RADIUS Username Password Credential Validator on page 912
- Configuring the Simple Username Password Credential Validator on page 914

Configuring the LDAP Username Password Credential Validator

Customize your LDAP Username Password Credential Validator's (PCV's) traits and behavior against your LDAP datastore to suit your needs.

About this task

The LDAP Username Password Credential Validator (PCV) verifies credentials using an organization's LDAP datastore.

When an authentication error occurs, PingFederate automatically parses the messages returned by PingDirectory, Microsoft Active Directory (AD), Oracle Unified Directory (OUD), or Oracle Directory Server (ODS) and categorizes them with error conditions.

When validating against a directory server other than PingDirectory, AD, OUD, or ODS, administrators can define custom message categorization by mapping specific error messages with wildcard support to the desired error conditions.

The error messages are returned to the HTML Form Adapter instances and the OAuth clients using the Resource Owner Password Credential grant type. The HTML Form Adapter is designed to show the error message it receives from the LDAP Username PCV. OAuth-client developers can create custom experiences based on the error responses, which contain the error messages. The HTML Form Adapter uses the relevant error conditions to determine the LDAP password-change scenarios and to present the relevant messages to the end users.



Z Tip:

These customizable messages are stored in the PingFederate message file, pingfederatemessages.properties, located in the <pf install>/pingfederate/server/default/conf/ language-packs directory.

You can localize these messages by using the PingFederate localization framework for an international audience. For more information, see .

Steps

- 1. Go to the **Instance Configuration** tab.
- 2. Optional: Override the authentication error messages.



You might require this option in order for a directory server other than PingDirectory, AD, OUD, or ODS to support the password change function in the HTML Form Adapter or to alter the end-user messages associated with that function.

- a. Click Add a new row to 'Authentication Error Overrides'.
- b. Enter an applicable LDAP error message under **Match Expression**.



You can use wildcard asterisks to match messages returned from your directory server. For example, *expired*.

- c. Select a relevant error condition from the **Error** list.
- d. Optional: Enter a key name or an error message under Message Properties Key.

No value

If you skip this field, PingFederate returns the default message based on the selected error condition.

A unique key name

If you enter a key name in this field and then add the key name with a key value (the desired error message) to the PingFederate message file, PingFederate returns that key value.

The key name must be unique. Furthermore, you may localize these messages by using the PingFederate localization framework for an international audience.

An error message

If you enter an error message in this field (without defining it in the PingFederate message file), PingFederate returns your message verbatim.

- e. Click **Update** under **Action**.
- f. Repeat these steps to add more overrides as needed.



Note:

Click Edit, Update, or Cancel to make or undo a change to an existing entry. Click Delete or Undelete to remove an existing entry or cancel the removal request.

Use the up and down arrows to change the display order. The display order does not affect runtime processing.

3. Select the LDAP datastore and enter information into the required fields. For more information about each field, see the following table.

Field	Description
LDAP Datastore	The LDAP datastore configured in PingFederate.
(Required)	If you have not yet configured the server to communicate with the directory server you need, click Manage Data Stores .
	Note: When connecting to an AD LDAP server, if you want to enable the
	password changes, password reset, or account unlock features in the HTML Form Adapter, you must secure the datastore connection to your AD LDAP server using LDAPS. AD requires this level of security to allow password changes.
Search Base	The location in the directory server from which the search begins.
(Required)	

Field Description Search Filter The LDAP query to locate a user record. (Required) If your use case requires the flexibility of allowing users to identify themselves using different attributes, you can include these attributes in your query. For instance, the following search filter allows users to sign on using either the samaccountName or employeeNumber attribute value through the HTML Form Adapter. (|(sAMAccountName=\${username})(employeeNumber= \${username})) **Important:** To ensure that your service providers (SPs) always get the expected attribute, select a specific user attribute as the source of the subject identifier when configuring the applicable SP connections. There are several ways to do so: Extend the PCV contract and fulfill the subject identifier through the HTML Form Adapter. For more information, see Extending the contract for the credential validator on page 914. Add a data source in the SP connection and fulfill the subject identifier through a datastore query. For more information, see Configuring attribute sources and user lookup on page 425. If you use authentication policy in conjunction with a policy contract, you can add a data source in the contract mapping configuration and fulfill the subject identifier in an SP connection through the authentication policy contract. For more information, see Applying policy contracts or identity profiles to authentication policies on page 250. When configuring multifactor authentication using PingID, where you chain an instance of the PingID Adapter behind an HTML Form Adapter instance, ensure that you also select a specific user attribute as the incoming user attribute for the PingID Adapter instance. For example, if you have set up PingFederate as the identity bridge for your PingOne for Enterprise account and have selected samaccountName as the subject identifier in the SP connection, you should also select samaccountName as the incoming user attribute for your PingID Adapter instance. You can accomplish this through an instance of the Composite Adapter or an authentication policy. For more information, see Input User ID Mapping in Configuring a Composite Adapter instance on page 283 or Incoming User ID in Specifying an incoming user ID on page 245, respectively. Scope of Search The level of search to perform in the search base. One Level indicates a search of objects immediately subordinate to the base object, not including the base object itself. Subtree indicates a search of the base object and the entire subtree within the base object distinguished name.

The deault selection is **Subtree**.

Field	Description
Field	Description
Case-Sensitive Matching	The option to enable case-sensitive matching between the LDAP error messages returned from the directory server and the Match Expression values specified on this window.
	This check box is selected by default.
Advanced fields for self-service password reset, account unlock, and user name recovery throu the HTML Form Adapter	
Display Name	The LDAP attribute used for personalizing messages to the users.
Attribute	This field is applicable for all password reset types (other than None), account unlock, and user name recovery.
	The default value is displayName.
Mail Attribute	The LDAP attribute containing the email address of the users.
(for password reset)	This field is required when password reset using one-time link or one-time password is enabled in any HTML Form Adapter instances that validate credentials against this LDAP Username PCV instance.
	Note:
	When configuring in conjunction with user name recovery, this attribute should correspond to the attribute specified on the left side of the Mail Search Filter field.
	The default value is mail.
SMS Attribute	The LDAP attribute containing the telephone number of the users.
(for password reset)	This field is required when password reset using text message is enabled in any HTML Form Adapter instances that validate credentials against this LDAP Username PCV instance.
	This field has no default value.
PingID Username	The LDAP attribute containing the PingID user name of the users.
Attribute	This field is required when password reset using PingID is enabled in any
(for password reset)	HTML Form Adapter instances that validate credentials against this LDAP Username PCV instance.
Mail Search Filter	The LDAP query to locate a user record using an email address, such as mail=\${mail}.
(for user name recovery)	This field is required when user name recovery is enabled in any HTML Form Adapter instances that validate credentials against this LDAP Username PCV instance.
	Note:
	When configuring in conjunction with password reset, the attribute specified on the left side of this search filter should correspond to the attribute specified in the Mail Attribute field.

Field	Description
Username Attribute	The LDAP attribute containing the user identifier of the users.
(for user name recovery)	This field is required when user name recovery is enabled in any HTML Form Adapter instances that validate credentials against this LDAP Username PCV instance.
	Note: This attribute should correspond to the attribute specified on the left side of the Search Filter field.
Mail Verified Attribute (for user name recovery)	The LDAP attribute indicating whether the user's email address is verified. The expected value of this user attribute is either true or false (case insensitive).
	This field is required when user name recovery using only verified email addresses is enabled in any HTML Form Adapter instances that validate credentials against this LDAP Username PCV instance.

Configuring the PingOne for Enterprise Directory Password Credential Validator

Use the PingOne for Enterprise Directory Username Password Credential Validator to verify credentials stored in your PingOne for Enterprise Directory.

Before you begin

To use the PingOne for Enterprise Directory Password Credential Validator, you must have:

- A PingOne for Enterprise account
- A PingFederate account

For more information, see *Managing PingOne directory users* in the PingOne for Enterprise Administration Guide.

Steps

On the **Instance Configuration** tab, enter your account information in **Client ID** and **Client Secret**. For more information about each field, refer to the following table. All fields are required.

Field	Description
Client ID	The REST API client ID is a unique identifier PingFederate uses to identify itself to the PingOne for Enterprise Directory API.
	For more information, see <i>View or renew directory API credentials</i> in the PingOne for Enterprise Administration Guide.
Client Secret	The client secret is used to authenticate the client ID against the PingOne for Enterprise Directory API.
	For more information, see <i>View or renew directory API credentials</i> in the PingOne for Enterprise Administration Guide.
Advanced Fields	
PingOne URL	The PingOne for Enterprise Directory API.
	The default value is https://directory-api.pingone.com/api.

Field	Description
Authenticate by Subject URL	The relative path for user authentication.
	The default value is /directory/users/authenticate?by=subject.
Reset Password URL	The relative path for password reset.
	The default value is /directory/users/password-reset.
SCIM User URL	The relative path for searching users requesting password reset.
	The default value is /directory/user.
Connection Pool Size	The maximum size of the connection pool to PingOne for Enterprise Directory.
	The default value is 100.
Connection Pool Idle Timeout	The maximum time (in milliseconds) that a connection can remain idle before it is closed and removed from the connection pool.
	The default value is 4000.

Configuring the RADIUS Username Password Credential Validator

Configure RADIUS servers to meet your authentication needs.

About this task

The RADIUS Username Password Credential Validator verifies credentials using the RADIUS protocol.

RADIUS supports strong authentication with both one-step (a combination of regular password and a onetime password in one field) and two-step (challenge-response) authentication. Two-step authentication is supported in the HTML Form Adapter.



Important:

If your RADIUS server is a Microsoft Network Policy Server (NPS), passwords containing special characters will not be encoded and decoded properly due to limitations with NPS.



RADIUS server messages are used by the HTML Form Adapter to determine the two-step authentication scenarios and to present a sign on window to the end users.

- 1. On the **Instance Configuration** tab, configure one or more RADIUS servers.
 - a. Click Add a new row to 'RADIUS Servers'.
 - b. In each field, enter the required information.

For more information about each field, refer to the following table. All fields are required.

Field	Description
Hostname	The IP address of the RADIUS server.
	For failover, enter one or more backup RADIUS servers by adding each server in its own row of the table. Each row represents a distinct RADIUS server that can be used for failover. PingFederate attempts to make a connection to each server in the order listed until a successful connection is obtained.
Authentication Port	The UDP port used to authenticate to the RADIUS server.
	The default value is 1812.
Authentication	The protocol used to authenticate to the RADIUS server.
Protocol	The available choices are Password Authentication Protocol (PAP) and Challenge Handshake Authentication Protocol (CHAP). Select the protocol expected by your RADIUS server.
	The default selection is PAP .
Shared Secret	The password shared between PingFederate and the RADIUS server used to encrypt the attribute identifying the NAS (Network Access Server) originating the request for access.



Note:

The NAS-IP-Address attribute is added to all Access-Request packets sent to the RADIUS server. The value is copied from the pf.engine.bind.address property in the <pf install>/ pingfederate/bin/run.properties file. Only IPv4 addresses are supported.

- c. Click **Update** in the **Action** column.
- d. Repeat these steps to add more RADIUS servers as needed.



Note:

Click Edit, Update, or Cancel to make or undo a change to an existing entry. Click Delete or **Undelete** to remove an existing entry or cancel the removal request.

Use the up and down arrows to adjust the order in which you want PingFederate to attempt credential authentication. If an earlier RADIUS server fails to validate the credentials, PingFederate moves sequentially through the list until credential validation succeeds. If none of the RADIUS servers is able to authenticate the user's credentials, the credential validation process fails.

2. Optional: Click **Show Advanced Fields** to reconfigure default settings.

For more information about each field, refer to the following table. All fields are required.

Field	Description
NAS Identifier	The password shared between PingFederate and the RADIUS server used to encrypt the attribute identifying the NAS (Network Access Server) originating the request for access.
	The default value is PingFederate.
Timeout	The maximum number of milliseconds before a connection timeout to the RADIUS server.
	The default value is 3000.
Retry Count	The number of times to retry a failed connection before moving to the next host.
	The default value is 3.

Configuring the Simple Username Password Credential Validator

The Simple Username Password Credential Validator verifies credentials maintained by PingFederate. This validator is best used for testing purposes or for an organization with few accounts.

Steps

- 1. On the Instance Configuration tab, click Add a new row to 'Users'.
- 2. Enter a user name, followed by a password (twice).
- Click Update in the Action column.
- 4. Repeat these steps to add more user credentials as needed.



Note:

Click Edit, Update, or Cancel to make or undo a change to an existing entry. Click Delete or **Undelete** to remove an existing entry or cancel the removal request.

Use the up and down arrows to adjust the order in which you want PingFederate to attempt credential authentication. PingFederate moves sequentially through the list until credential validation succeeds or no match is found.

Extending the contract for the credential validator

Extend Password Credential Validator (PCV) instance contracts to return attribute values relevant to authenticated users.

About this task

In some use cases, you might want to extend the contracts of the PCV instance. For example, you might use extended attributes to map into a USER KEY for an OAuth persistent grant configuration.

This capability allows the validator to return attribute values pertaining to the authenticated users from PingOne for Enterprise Directory, a directory server, or a RADIUS server.



If you are configuring an HTML Form Adapter instance with an instance of the LDAP Username Password Credential Validator, extend the contract of the adapter by the same attribute names in order for the credential validator to pass extended attribute values to the HTML Form Adapter instance.

If you are configuring the HTML Form Adapter instance with an instance of the RADIUS Username Password Credential Validator, you only need to extend the contract of the HTML Form Adapter instance itself.

Steps

1. Copy the vendor-specific attribute dictionaries into the pingfederate/server/default/conf/ radius directory.



Note:

The format of the dictionaries must use the *FreeRadius dictionary syntax*.

- 2. Edit the existing dictionary file to include each of the dictonaries.
- 3. Optional: On the Extended Contract tab, enter an attribute name and click Add.



Note:

Click Edit, Update, or Cancel to make or undo a change to an existing entry. Click Delete or **Undelete** to remove an existing entry or cancel the removal request.

Finishing the Password Credential Validator instance configuration

On the **Summary** tab, review your configuration to determine whether to keep it, change it, or remove it entirely.

Steps

- To keep your changes, click Save.
- To amend your configuration, click the name of the corresponding tab and then follow the configuration. wizard to complete the task.
- To discard your changes, click Cancel.

Active Directory and Kerberos

You can configure PingFederate to authenticate users through the following identity provider (IdP) adapters or token processors.

Adapter or Token Processor	Description	
PingFederate integrated Kerberos Adapter	Using the built-in Kerberos Adapter with a configured AD domain allows a PingFederate identity provider (IdP) server to perform single sign-on (SSO) to service provider (SP) applications based on Kerberos tickets.	
PingFederate integrated Kerberos Token Processor	The built-in Kerberos Token Processor accepts and validates Kerberos tokens through a configured Kerberos Realm from a web service client.	
Integrated Windows Authentication (IWA) Integration Kit (version 3.0 and later)	Using the separately available IWA Adapter with a configured AD domain allows a PingFederate IdP server to perform SSO to SP applications based on IWA credentials.	

Configuring Active Directory domains or Kerberos realms

You can configure an Active Directory (AD) domain or Kerberos realm to authenticate users.

Steps

- 1. Go to System # Data & Credential Stores # Active Directory Domains/Kerberos Realms.
- 2. From the Manage AD Domains/Kerberos Realms window, configure the AD environment to integrate with PingFederate. For more information, see Configuring the Active Directory environment on page 916.
- 3. Click Add Domain/Realm to create an AD domain.



Important:

Do not configure subdomains if the parent domain in the same forest is already configured. For more information, see *Multiple-domain support* on page 916.



Note:

Click the name of an existing domain to edit it. Use the **Delete** and **Undelete** links to remove a domain or cancel a removal request.

Multiple-domain support

If your network uses multiple domains in a single server forest, configure one domain within PingFederate if there is a trust relationship with the other domains you want to use.

This configuration requires a trust relationship among domains, which is established by default when subdomains or separate domains are created within the same forest. For more information, see https:// docs.microsoft.com/en-us/previous-versions/windows/it-pro/windows-server-2003/cc773178(v=ws.10)? redirectedfrom=MSDN.



Note:

If you are configuring only one domain, then you also need to configure only one Service Principal Name. For more information, see Configuring the Active Directory environment on page 916.

If your network topology consists of multiple forests without a trust relationship between them, you must configure multiple adapter or token processor instances. Map each instance to a separate domain and then map these adapter or token processor instances to your service provider (SP) connections that authenticate using the integrated Kerberos Adapter, the integrated Kerberos Token Processor, or the separately available IWA Adapter.

For information about configuring the PingFederate Integrated Windows Authentication (IWA) adapter for multiple-domain Active Directory trusts, see https://support.pingidentity.com/s/article/How-to-configure-IWA-with-multiple-Active-Directory-trusts.

Configuring the Active Directory environment

Configure Active Directory to access a domain and enable Kerberos as an authentication option for it.

About this task

To enable Kerberos authentication, you must make several Active Directory configuration changes to grant PingFederate access to the domain and add the domain to PingFederate.



Important:

Do not configure subdomains if the parent domain in the same forest is already configured. For more information, see *Multiple-domain support* on page 916.



Note:

You must have Domain Administrator permissions to make the required changes.

Steps

- 1. Create a domain user account that PingFederate can use to contact the Kerberos Key Distribution Center (KDC). The account should belong to the Domain Users group. We recommend that you set the password with no expiration.
- 2. Use the Windows utility setspn to register Service Principal Name (SPN) directory properties for the account by executing the following command on the domain controller.

setspn -s HTTP/<pf-idp.domain.name> <pf-server-account-name> , where <pfidp.domain.name> is the canonical name of the PingFederate server and <pf-server-account-name> is the domain account you want to use for Kerberos authentication. For more information on canonical name, see https://tools.ietf.org/html/rfc2181#section-10.



When executing the setspn command, you must capitalize HTTP and follow it with a forward slash (/).

3. Verify that the registration was successful by executing the following command.

```
setspn -l cpf-server-account-name>
```

This gives you a list of SPNs for the account. Verify that HTTP/<pf-idp.domain.name> is one of them.



Note:

After making an SPN change, any authenticated end users must re-authenticate by closing the browser or signing off and back on before attempting single sign-on (SSO).

Adding a domain

Configure Active Directory domains or Kerberos realms that PingFederate can use to contact the domain controllers or the key distribution centers (KDCs) for verifying user authentication.

Steps

In the Manage Domain/Realm window, enter the required information based on the following table.

Field	Description
Domain/Realm Name	The fully-qualified domain or realm name. For example, companydomain.com
Domain/Realm Username	The ID for the domain or realm account name.
Domain/Realm Password	The password for the domain or realm account.

Field	Description
Domain Controller/Key Distribution Center	Specify the host name or IP address of your domain controller or KDC, such as dc01-yvr, and then click Add . Repeat this step to add multiple servers.
Host Names (optional)	If a host name is used, PingFederate appends the domain to the host name to formulate the fully qualified domain name (FQDN) of the server unless the Suppress DC / Domain Concatenation check box is selected.
	If unspecified, PingFederate uses a DNS lookup.
Suppress DC / Domain Concatenation	Select this check box to specify the desired FQDNs under Domain Controller/ Key Distribution Center Host Names . When selected, PingFederate does not append the domain to the host names.
Test Domain/Realm Connectivity	Tests access to the domain controller or KDC from the administrative-console server.
	When a connection to any of the configured controllers or KDCs is successful, the message <code>Test Successful</code> appears. Otherwise, the test returns error messages near the top of the window.
	Tip:
	For help resolving connectivity issues, select the Debug Log Output check box on the Manage Domain/Realm Settings window, run the test again, and review the debug messages in the PingFederate server log.
	This test stops at the first successful result when multiple domain controllers or KDCs are specified, so not all servers are necessarily verified. Depending on the network architecture, the engine nodes deployed in a cluster might establish connections differently. As a result, the engine nodes and the console node might connect to different domain controllers or KDCs.

Managing domain connectivity settings

You can change the default security and logging settings for all configured Active Directory domains and Kerberos realms.

Steps

• On the **Manage Domain/Realm Settings** tab, change the default transport protocol, the debug option, the timeout value, and the number of retry attempts. For more information, refer to the following table.

Field	Description
Force TCP	When selected, requires use of the Transmission Control Protocol instead of the default User Datagram Protocol. Use this option when firewall or network configurations require acknowledgment that packets are properly received.
	Note: If you choose this option, you must restart PingFederate after saving the configuration.
Debug Log Output	When selected, sends verbose messages to the PingFederate server log for all interactions with the domain controllers or the Key Distribution Centers (KDCs).

Field	Description
AD Domain Controller/Key Distribution Center Timeout (secs)	Specifies the amount of time (in seconds) PingFederate waits for a network response from a domain controller or KDC. The default is 3.
	Note:
	This value applies to each attempt PingFederate makes to contact the domain controller or KDC.
	The new timeout takes effect only after you save the configuration and restart PingFederate
AD Domain Controller/Key Distribution Center Retries	Specifies the number of times PingFederate tries contacting the domain controller or KDC. The default is 3.

External systems

Many use cases require communications between PingFederate and other systems, such as PingOne or a database server. System # External Systems is where you can set up and maintain such integrations.

Connecting to PingOne for Enterprise after initial setup

Integrating PingOne for Enterprise with PingFederate provides a powerful solution combining the benefits of an on-premise deployment with the flexibility of a cloud solution.

Steps

- 1. Go to System # External Systems # Connect to PingOne for Enterprise.
- 2. On the Connect to PingOne for Enterprise window, click Sign on to PingOne to get your activation key.
- 3. Sign on using your PingOne Admin Portal credentials.



If you do not have an account, you can register for a free trial.

- 4. Copy the **Activation Key** value from the PingOne admin portal.
- 5. Close the browser tab and go back to the PingFederate administrative console.
- 6. On the PingOne Account window, paste the key value in the Activation Key field. Click Next.

Configuring identity repository settings

Set up a customized directory configure identity repository settings either immediately or at a later time.

About this task

On the Identities tab, you can optionally connect to a directory server.

Steps

- Go to System # External Systems # Connect to PingOne for Enterprise and access the Identities tab.

To enable directory integration, select Yes, Connect a Directory Server.

You can create a new datastore or reuse an existing datastore in this configuration.

Create a new datastore

Provide the required information to connect to a directory server, and then click **Next**. More information about each field is provided in the following table.

Field	Description
Directory Type	Select the type of directory server from the list.
	See <i>System requirements</i> on page 110 for a list of supported directory servers.
Data Store Name	Enter the name of the datastore.
Hostname	Enter the location of the directory server.
	It can be the IP address, the host name, or the fully qualified domain name of the directory server. The entry might include a port number.
Service Account DN	Enter the distinguished name (DN) of the service account that PingFederate can use to communicate with the directory server.
Password	Enter the password associated with the service account.
Search Base	Enter the DN of the location in the directory where PingFederate begins its datastore queries.
Search Filter	Enter the LDAP query to locate a user record for attribute lookup and potentially credential validation.
	The default value is either sAMAccountName=\${username} or uid=\${username}, depending on the selected directory type.
	If you require a more advanced search filter, ensure the value is a valid LDAP filter. For more information, consult your directory administrators.

When you click Next, PingFederate tries to establish a secure (LDAPS) connection to the directory server.

If the directory server does not support LDAPS, the Unsecure Connection window appears. If you want to continue without a secure connection, click Next. Alternatively, you can go back to the **Identities** tab and specify a different directory server.

If the certificate presented by the directory server is not trusted by PingFederate, the Certificate Error window appears. You can import the certificate used by the directory server to establish a secure connection, and then click Next in the Identities tab and specify a different directory server.

Use an existing datastore

Click Begin, and then follow the on-screen instructions to create a service provider (SP) connection to PingOne for Enterprise.

 Optional: To set up a directory later, select No, Don't Connect a Directory Server and then click Next.



This setup scenario is suitable for proof of concept. Multiple local test accounts are created as a result.

Configure PingOne for Enterprise's general settings and single sign-on (SSO) settings, and enable and configure a RADIUS server to integrate PingID with a VPN.

About this task

Go to **System # External Systems # PingOne for Enterprise Settings**, configure various PingOne for Enterprise integration settings and optionally enable and configure a built-in RADIUS server to integrate PingID with your VPN.

Configuring PingOne for Enterprise settings

To configure the PingOne for Enterprise settings, adjust the various integration settings.

Steps

- 1. Go to System # PingOne for Enterprise Settings.
- 2. Configure the PingOne for Enterprise integration settings. For more information, see the following table.

Field	Description
Enable Single Sign-On from PingOne to the PingFederate Administrative Console	Toggles the ability to sign on to the administrative console using the PingOne admin portal credentials.
Enable Monitoring of PingFederate from PingOne	Toggles the ability to monitor your PingFederate server (or servers in a clustered environment) from the PingOne admin portal
Rotate Key	Update the authentication key that PingFederate uses to communicate with PingOne for Enterprise.
	Note: Periodic rotation can ensure optimal security of your environment.
	Note: PingFederate also automatically rotates the signing certificate used by the managed service provider (SP) connection. For more information, see <i>Managed SP connection to PingOne for Enterprise and signing certificate</i> on page 613.
Launch PingOne Admin Portal	Use to access the PingOne admin portal.

Field	Description
Disconnect from PingOne	Use to disconnect PingFederate from your PingOne account.
	This is applicable if you have made changes that you should not propagate to your PingOne for Enterprise account.
	For instance, you have two PingFederate environments, testing and production. The production PingFederate server is configured with a managed SP connection to PingOne for Enterprise, but the PingFederate test server is not. You have just exported a configuration archive from the production server and imported it to the test server. As soon as the configuration archive is imported, the administrative console prompts you to decide whether to update PingOne for Enterprise or to disconnect from PingOne for Enterprise. In this example, you should disconnect the test server from PingOne for Enterprise so that nothing is uploaded to your PingOne for Enterprise account from the test server.

3. Save your configuration.

Configuring PingOne SSO settings

About this task

Use PingOne for Enterprise to enable single sign-on (SSO) and upload configuration changes to your account.

Steps

 To enable single sign-on (SSO) through PingOne for Enterprise, click Identity Repository Configuration.



Note:

This is applicable if you have not yet completed the PingOne SSO configuration in the past, which would have created a managed service provider (SP) connection to PingOne for Enterprise.

To upload configuration changes to your PingOne account, go to the PingOne for Enterprise settings window and select Update PingOne Identity Repository, then confirm your decision.

This is applicable if you have made changes that you should propagate to your PingOne for Enterprise account.

For example, you are about to set up a new SAML application on PingOne for Enterprise that requires a telephone number of the user. Because the current attribute contract in the managed SP connection does not include an attribute for telephone number, you extend the attribute with a new attribute, PrimaryTelephone. After the connection is saved, the administrative console prompts you to decide whether to update PingOne for Enterprise or to disconnect from PingOne for Enterprise. In this example, you should upload the new configuration to PingOne for Enterprise so that the new PrimaryTelephone attribute is made available when you set up the new SAML application in PingOne for Enterprise.

Enabling and configuring the built-in RADIUS server to integrate PingID with your VPN

About this task

Use the PingID VPN (RADIUS) configuration wizard to merge PingID with your own VPN.

- 1. Click PingID Configuration to open the PingID VPN (RADIUS) configuration wizard.
 - This is applicable if you have not completed the **PingID VPN (RADIUS)** configuration in the past, which would have created a PingID provisioning connection, an instance of the PingID password credential validator (PCV), or both.
- 2. Use the wizard to configure the server.

Configuring SSO from PingOne admin portal to PingFederate administrative console

You can single sign-on (SSO) to the PingFederate administrative console from PingOne for Enterprise and configure authentication procedures as desired.

About this task

In PingFederate 10.1 and later, you can connect to PingOne for Enterprise after the initial PingFederate setup by going to **System # External Systems # Connect to PingOne for Enterprise**.

Additionally, you can continue to sign on to the administrative console through native or alternative console authentication using the direct sign on page. You can also disable the direct sign on page to enforce the policy that administrators must SSO to the administrative console from the PingOne admin portal.

Steps

- To SSO to the administrative console:
 - a. Start a web browser.
 - b. Browse to the URL https://<pf_host>:9999/pingfederate/app, where <pf_host> is the network address of your PingFederate server, either an IP address, a host name, or a fully qualified domain name reachable from your computer.

Result:

If the SSO option is enabled on the **PingOne for Enterprise Settings** window and you have signed on to the PingOne admin portal, the PingFederate administrative console is made available. If you are not signed on to the PingOne admin portal, you are prompted to enter your PingOne admin portal credentials. Upon verification, the PingFederate administrative console is made available.

- To sign on through native or alternative console authentication:
 - a. Start a web browser.
 - b. Browse to the URL https://*<pf_host>*:9999/pingfederate/app?service=page/directLogin, where *<pf_host>* is the network address of your PingFederate server, either an IP address, a host name, or a fully qualified domain name reachable from your computer.
- To disable native and alternative console authentication:
 - a. Edit the <pf install>/pingfederate/bin/run.properties file.
 - b. Change the pf.console.authentication property value to none.
 - c. Save the change and then restart PingFederate.



Note:

In a clustered PingFederate environment, you only need to modify the run.properties file on the console node.

Result:

After restart, the direct login page is disabled. Administrators can only SSO to the PingFederate administrative console from the PingOne admin portal at https://cpf_host>:9999/pingfederate/app.

To re-enable native or alternative console authentication, update the pf.console.authentication property accordingly and then restart PingFederate.

Monitoring PingFederate from the PingOne admin portal

After connecting PingFederate to PingOne for Enterprise, you can monitor PingFederate from the PingOne admin portal.

About this task

The PingOne admin portal displays your PingFederate server (or servers if you have a clustered PingFederate environment) with basic information such as the node index number, the IP address, and the connection status with the date last seen. For each server, you can also collect additional information such as CPU load, Java virtual machine (JVM) memory information, and system memory information.

Steps

- 1. Go to System # External Systems # PingOne for Enterprise Settings.
- 2. Optional: If you do not want to monitor PingFederate using the PingOne admin portal, clear the **Enable** Monitoring of PingFederate From PingOne check box and click Save.

Updating the PingOne identity repository

You can upload configuration changes to your connected PingOne for Enterprise account's admin portal or disconnect from your PingOne for Enterprise account if needed.

About this task

After a managed service provider (SP) connection to PingOne Admin Portal is established, PingFederate monitors configuration changes that can impact the connection, such as updates to the base URL or imports of configuration archives that include managed SP connections to PingOne for Enterprise. When PingFederate detects such changes, the administrative console prompts you to decide whether to update PingOne for Enterprise or to disconnect from PingOne for Enterprise in a banner message.

Steps

Go to System # External Systems # PingOne for Enterprise settings. Choose from:

 To upload configuration changes to your PingOne for Enterprise account, click Update PingOne Identity Repository.



Note:

This is applicable if you have made changes that you should propagate to your PingOne for Enterprise account.

For example, you are about to set up a new SAML application on PingOne for Enterprise that requires a telephone number of the user. Because the current attribute contract in the managed SP connection does not include an attribute for telephone number, you extend the attribute with a new attribute, PrimaryTelephone. After the connection is saved, the administrative console prompts you to decide whether to update PingOne for Enterprise or to disconnect from PingOne for Enterprise. In this example, you should upload the new configuration to PingOne for Enterprise so that the new PrimaryTelephone attribute is made available when you set up the new SAML application in PingOne for Enterprise.

To disconnect PingFederate from your PingOne for Enterprise account, click Disconnect from PingOne.



Note:

This is applicable if you have made changes that you should not propagate to your PingOne for Enterprise account.

For instance, you have two PingFederate environments, testing and production. The production PingFederate server is configured with a managed SP connection to PingOne for Enterprise, but the PingFederate test server is not. You have just exported a configuration archive from the production server and imported it to the test server. As soon as the configuration archive is imported, the administrative console prompts you to decide whether to update PingOne for Enterprise or to disconnect from PingOne for Enterprise. In this example, you should disconnect the test server from PingOne for Enterprise so that nothing is uploaded to your PingOne for Enterprise account from the test server.

Managing CAPTCHA settings

Use PingFederate's functionality to enter and install Google CAPTCHA.

About this task

PingFederate supports invisible reCAPTCHA from Google.

Steps

- 1. Go to System # External Systems # CAPTCHA Settings, and enter the site key assigned to your account by Google.
- 2. Enter the associated secret key.
- 3. Click Save.

Managing SMS provider settings

To connect PingFederate to Twilio as an SMS provider through which PingFederate can send text message notifications for self-service password reset requests, enter the required information based on your Twilio account.

Steps

- 1. Go to System # External Systems # SMS Provider Settings.
- 2. Click Manage SMS Provider Settings.
- 3. On the **SMS Provider Settings** window, enter the required information.

Field	Description
Account SID	The account number assigned to your account by Twilio.
Auth Token	The password assigned to your account by Twilio. Used in conjunction with the account number to authenticate with Twilio when PingFederate makes outbound API calls for the purpose of sending text message notifications to the intended recipients.
From Number	The sender number in the text message notifications.



For additional information about each field or Twilio, see http://support.twilio.com.

4. Click Save.

Result

After you have saved them, these SMS provider settings apply to all services using text message notifications.

Managing notification publisher instances

Use PingFederate's functionality to create, edit, review, delete, or set as a default any notification publisher instance.

About this task

PingFederate delivers messages to administrators and end users based on notification publisher settings. Depending on your use cases, you can create one or more notification publisher instances. For example, you might select an SMTP notification publisher instance to deliver messages to your end users in an HTML Form Adapter instance, another SMTP notification publisher instance to deliver licensing messages to your fellow administrators, and an Amazon SNS notification publisher instance to deliver messages regarding SAML metadata updates.

Steps

- Go to System # External Systems # Notification Publishers.
 Choose from:
 - To configure a new instance, click Create New Instance.
 - To modify an existing instance, select it by its name under Instance Name.
 - To review the usage of an existing instance, click Check Usage under Action.
 - To remove an existing instance or to cancel the removal request, click **Delete** or **Undelete** under **Action**.
 - To make an instance the default notification publisher instance, click Set as Default under Action.

Defining a notification publisher instance

Define a notification publisher instance.

Steps

- 1. Go to System # External Systems # Notification Publishers.
- 2. Click Create New Instance.
- 3. On the **Type** tab, enter an instance name and an instance ID.
- 4. From the **Type** list, select the type of the notification publisher instance that you want to use.
- 5. Click Next.

Notification publisher instance configurations

You can configure a notification publisher instance based on its type.

This configuration varies depending on the type of the notification publisher.

For specific configuration steps, see the following topics..

- Configuring an Amazon SNS Notification Publisher instance on page 927
- Configuring an SMTP Notification Publisher instance on page 931

Create and configure an Amazon SNS notification publisher instance to help subscribers effectively direct their messages to their intended audiences.

About this task

When using an Amazon SNS notification publisher, PingFederate is the publisher, and the intended recipients are the subscribers. Topics are the destinations to which PingFederate publishes messages. When configuring an Amazon SNS notification publisher instance, you must specify an Amazon SNS topic.

For more information about Amazon SNS and topic management, see the AWS documentation on https://docs.aws.amazon.com/sns/latest/dg/welcome.html.

Steps

1. On the **Instance Configuration** tab, configure the notification publisher instance as follows.

Field	Description
SNS Topic ARN	The Amazon Resource Name (ARN) topic to which PingFederate publishes messages. Enter an ARN in the following format.
	<pre>arn:aws:[service]:[region]:[accountid]:[resourceType/ resourcePath]</pre>
Max Payload Size	The maximum payload size in kilobytes. Enter a value between 1 and 8192.
	Click Show Advanced Fields to reveal this field. The default value is 256.

2. Click **Next** to view the **Summary** tab, and then click **Save**.

Result

PingFederate categorizes notification messages into various event types. Each event type comes with a set of relevant information to help subscribers craft the final message for the intended audience.

Event types and variables

This lists the various event types and their respective keys (variables) used within the PingFederate Amazon SNS notification publisher instance configuration.

Message payload

As a publisher, PingFederate creates notification messages in JSON format and sends them to the configured topic. This JSON message body contains two top-level keys: data and configuration, as illustrated in the following snippet.

```
"data": {
    "USERNAME": "jdoe",
    ...
},
"configuration": {
    "com.pingidentity.notification.config.locale": "en-US",
    ...
    "com.pingidentity.notification.config.event.type":
"ADMIN_PASSWORD_CHANGED"
}
```

For all events, PingFederate provides relevant information by including various key: value pairs in the message body, located inside the value of the data key.

The value of the com.pingidentity.notification.config.event.type key, located inside the value of the configuration key, indicates the event type. In this example, the event type is ADMIN PASSWORD CHANGED.

For end user-oriented events, the value of the com.pingidentity.notification.config.locale key, also located inside the value of the configuration key, indicates the locale of the end user who initiates the request.

Review the following sections for more information on event types and their respective keys, which are referred to as variables.

Events for administrators

Local administrative account management events

Event type Variables

ADMIN_ACCOUNT_CHANGE_NOTIFICATION OF Fesents the username of the local administrative account who has turned off the Notify Administrator of Account Changes option.

- RECEIVER represents the email addresses of all the local administrative accounts configured with an email address.
- NOTIFY represents the Notify Administrator of Account Change option on the Administrative Accounts window.
- CURRENT USER MESSAGE represents the username of the administrator who initiated the change.



Unless otherwise noted, the rest of the variables in the Administrative Accounts section are either self-explanatory or identical to those mentioned here.

ADMIN_EMAIL_CHANGED

- USERNAME
- RECEIVER
- DEPARTMENT
- DESCRIPTION
- PHONE NUMBER
- CURRENT USER MESSAGE

PingFederate sends two messages for this event type. Variables and their values remain the same, except for the RECEIVER value. They are intended to notify the end user at both the previous email address and the new email address.

ADMIN_PASSWORD_CHANGED USERNAME

- RECEIVER
- DEPARTMENT
- DESCRIPTION
- PHONE NUMBER
- CURRENT USER MESSAGE

Variables Event type CERTIFICATE_EVENT_ACTIVATEDSERIAL NUMBER SUBJECT DN CERTIFICATE_EVENT_CREATED EX DATE PENDING CERT SERIAL NUM PENDING EX DATE • ACTIVE CONNECTIONS represents the connections impacted by the creation of the pending certificate and the activation of it. ACTIVATION DATE CERTIFICATE_EVENT_EXPIRED.CERTIFICATELENT_FINAL_WARN, CERTIFICATE_EVENT_INITIAL_WARN DATE • EX TYPE • CONN NAME represents the connection impacted by any of the three certificate expiration events. DAYS LEFT ACTION $SAML_METADATA_UPDATE_EVENENENENTYTY_DID_NOT_FOUND$ - CONNECTION NAME METADATA URL • METADATA URL NAME SAML_METADATA_UPDATE_EVENMETANATA URL METADATA URL NAME SAML_METADATA_UPDATE_EVENTNUPPPATED - CONNECTION NAME • METADATA URL UPDATED represents any updated connection settings. • OUT OF SYNC represents any out-of-sync connection settings. SERVER_LICENSING_EVENT_WARKINGATE SERVER_LICENSING_EVENT_EXPIRED, LEFT SERVER_LICENSING_EVENT_SHUTDOWN

Events for end users

Self-service password management, account recovery, and username recovery

Event type	Variables
ACCOUNT_UNLOCKED	 USERNAME represents the user name of the end user where the request is made. RECEIVER represents the email address of the end user where the request is made. ADAPTER_ID represents the Instance ID of the invoking HTML Form Adapter instance. PCV_ID represents the Instance ID of the Password Credential Validator (PCV) instance involved.
	Note: Unless otherwise noted, the rest of the variables in the HTML Form Adapter instances section are either self-explanatory or identical to those mentioned here.
PASSWORD_CHANGED	GIVEN_NAMEUSERNAMERECEIVERADAPTER_IDPCV_ID
PASSWORD_RESET	 USERNAME RECEIVER ADAPTER_ID PCV_ID STATUS
PASSWORD_RESET_FAILED	USERNAMERECEIVERADAPTER_IDPCV_ID
PASSWORD_RESET_ONE_TIN	
USERNAME_RECOVERY	USERNAMERECEIVERADAPTER_IDPCV_IDDISPLAY_NAME

Customer IAM email ownership verification

Event type	Variables
OWNERSHIP_VERIFICATION_	ONE USERVALLE represents the user name of the end user who should receive an email ownership verification request. RECEIVER represents the email address to which the email ownership verification request is sent. CODE represents the one-time hyperlink that the end user can use to verify the ownership of the email address associated with the account.

Configuring an SMTP Notification Publisher instance

Set up an instance of the SMTP Notification Publisher for PingFederate to notify administrators and end users about various events. You can configure multiple instances, each with different settings as needed.

Steps

1. On the **Instance Configuration** tab, provide the required information or update any default or previously configured setting values.

For more information about each field, refer to the following table.

Field	Description
From Address	The email address that appears in the "From" header line in email messages generated by PingFederate. The address needs a valid format
(Required)	but does not need to be set up on your system.
Email Server	The IP address or host name of your email server.
(Required)	
SMTP Port	The SMTP port on your email server.
	The default value is 25.
Encryption Method	Select SSL/TLS to establish a secure connection to the email server at the SMTPS port.
	Select STARTTLS to establish an unencrypted connection to the email server at the SMTP port and initiate a secure channel afterward.
	Select None , the default value, to establish an unencrypted connection to the email server at the SMTP port.
SMTPS Port	The secure SMTP port on your email server. Not applicable unless SSL/TLS is the chosen encryption method.
	The default value is 465.
Verify Hostname	Indicates whether to verify the host name of the email server matches the Subject (CN) or one of the Subject Alternative Names from the certificate.Not applicable unless either SSL/TLS or STARTTLS is the chosen encryption method.
	This check box is selected by default.

Field	Description
UTF-8 Message Header Support	Indicates whether the email server supports UTF-8 encoding in message headers; for example, in the recipient email address. Enable this option only if your email server supports this feature.
	With this option enabled, PingFederate supports UTF-8 characters in the sender and recipient email addresses. It does not support Emojis in the domain portion of the email address.
Username	Authorized email account.
Password	Password for the authorized email account.
Test Address	Enter an email address PingFederate should use to verify connectivity with the configured email server.
Click Show Advanced I	Fields to review the following settings. Modify as needed.
Connection Timeout	The amount of time in seconds that PingFederate waits before it times out connecting to the SMTP server.
	The default value is 30.
Retry Attempt	The number of times that PingFederate tries to connect to the SMTP server after encountering an error.
	The default value is 2.
Retry Delay	The amount of time in minutes that PingFederate waits before trying to send an email message again.
	The default value is 2.
Enable SMTP Debugging Messages	Turns on detailed error messages for the PingFederate server log to help troubleshoot SMTP issues.
	CAUTION:
	This setting is disabled by default. When enabled, PingFederate logs email messages, which can contain sensitive information, to the server log. Consider enabling debug messages solely for troubleshooting purposes and disabling this option when debug messages are no longer required.

2. Click Next.

Next steps

You can modify email-notification template files to suit the particular branding requirements. For more information, see *Customizable email notifications* on page 810.

Finalizing actions for a notification publisher instance

Finish any required tasks presented to you on the **Actions** tab to complete configuration for your notification publisher instance.

About this task



Note

Depending on the data store implementation, configuration requirements vary. If no action is required, the **Actions** tab is not shown.

Steps

• On the **Actions** tab, perform any required tasks, and click **Next**.

Reviewing a notification publisher instance configuration On the **Summary** tab, review your configuration.

Steps

Perform the following actions as needed.

Action	How to accomplish it
	Click the corresponding tab and follow the configuration workflow
Keep your changes	Click Done and continue with the configuration
Discard your changes	Click Cancel

Result

Once set up and saved, you can select this notification publisher instance in any component that is capable of triggering or handling events.

Troubleshooting

Basic troubleshooting tips are provided here to help you overcome common difficulties with PingFederate.

- Enabling debug messages and console logging
- Resolving startup issues
- Troubleshooting data store issues
- Resolving URL-related errors
- Resolving service-related errors
- Troubleshooting authentication policy issues
- Troubleshooting registration and profile management issues
- Troubleshooting runtime errors
- Troubleshooting OAuth transactions
- Other runtime issues
- Collecting support data

Help is also available from the Support Center.

Enabling debug messages and console logging

Make changes to the $< pf_install > / pingfederate/server/default/conf/log4j2.xml file to enable debug messages and console logging.$

About this task

For troubleshooting purposes, you can adjust the log level to DEBUG in the log4j2.xml file and re-enable console logging.



When you no longer require debug messages and console logging, turn them off.

Steps

- Edit the <pf_install>/pingfederate/server/default/conf/log4j2.xml file.
 - a. To enable verbose messages, look for Limit categories inside the Loggers element, particularly the five logger names shown in bold in the following snippet.

```
<Loggers>
    <!-- ========= -->
    <!-- Limit categories -->
    <!-- ========== -->
    <Logger name="org.sourceid" level="INFO" />
    <Logger name="org.sourceid.saml20.util.SystemUtil" level="INFO"</pre>
additivity="false">
        <AppenderRef ref="CONSOLE" />
        <AppenderRef ref="FILE" />
    </Logger>
    <Logger name="com.pingidentity" level="INFO" />
    <Logger name="com.pingidentity.common.util.ErrorHandler" level="INFO"</pre>
additivity="false">
        <AppenderRef ref="CONSOLE" />
        <AppenderRef ref="FILE" />
    </Logger>
    <Logger name="com.pingidentity.appserver.jetty.PingFederateInit"</pre>
level="ERROR" additivity="false" includeLocation="false">
        <AppenderRef ref="CONSOLE" />
    </Logger>
</Loggers>
```

Then update the five loggers as shown in bold in the following snippet.

```
</Logger>
    <Logger name="com.pingidentity.appserver.jetty.PingFederateInit"</pre>
level="INFO" additivity="false" includeLocation="false">
        <AppenderRef ref="CONSOLE" />
        <AppenderRef ref="FILE" />
    </Logger>
</Loggers>
```

Note:

In this snippet, console logging is enabled for loggers

```
org.sourceid.saml20.util.SystemUtil,
com.pingidentity.common.util.ErrorHandler, and
com.pingidentity.appserver.jetty.PingFederateInit.
```

Console logging can be resource intensive. If you do not require console logging or prefer to review the server log instead, comment out the <appenderRef ref="CONSOLE" /> entry for these three loggers.



As needed, you can update the log level for other loggers.

b. To enable console logging, look for the Set up the Root logger section as shown in the following example.

```
<!-- Set up the Root logger -->
<AsyncRoot level="INFO" includeLocation="false">
  <!-- <AppenderRef ref="CONSOLE" /> -->
  <AppenderRef ref="FILE" />
</AsyncRoot>
```

Then, update as shown in bold in the following snippet.

```
<!-- Set up the Root logger -->
<AsyncRoot level="INFO" includeLocation="false">
  <AppenderRef ref="CONSOLE" />
  <AppenderRef ref="FILE" />
</AsyncRoot>
```



Important:

When you no longer require console logging, comment out the AppenderRef ref="CONSOLE" /> entry for the AsyncRoot logger.

2. Save any changes made.

Result

Changes, such as adding a Logger or adjusting log levels, are activated within half a minute. You do not need to restart PingFederate.

Resolving startup issues

Resolve PingFederate startup issues related to Java runtime and the license file installs.

PingFederate startup problems and solutions

Problem	Solution
PingFederate does not start.	Make sure that a supported Java runtime is installed. For more information, see <i>Installing Java</i> on page 119.
The server starts but indicates the license file is not found or invalid.	Ensure a current license is installed. For more information, see <i>Reviewing license information</i> on page 848.

Troubleshooting data store issues

Resolve PingFederate data store issues related to failing to establish a connection to a data store and directory.

Data store connection problems and solutions

Problem	Solution
When setting up the JDBC data store, a connection cannot be established.	Verify that the proper drivers and connectors have been installed.
	Also, verify the connection URL, user name, and password. If unsuccessful, contact your database administrator.
	For more information, see Configuring a JDBC connection on page 868.
Cannot connect to a Directory Service with the LDAP protocol.	Verify the connection URL, port, principal, and credentials. If unsuccessful, contact your system administrator. For more information, see <i>Configuring an LDAP connection</i> on page 871.
	If using LDAPS, ensure the LDAP server's SSL certificate is signed by a trusted certificate authority or is imported into PingFederate. For more information, see <i>Manage trusted certificate authorities</i> on page 600.

Resolving URL-related errors

Review and update the URL of the identity provider (IdP) initiated single sign-on to resolve the 404 Not Found and System Error messages.

If a user encounters a 404 Not Found status or a System Error message, check the URL of the request.

Example

Examples

404 Not Found

https://sso.idp.local/idp/startSSO.ping&PartnerSpId=sp1&TargetResource=https%3A%2F %2Fapp.sp1.local%2F causes a 404 Not Found error, because the separator between the path of the URL and the first query parameter is incorrect. The correct separator is a question mark? and not an ampersand &.

System Error

https://sso.idp.local/idp/startSSO.ping?PartnerSpId=sp1?TargetResource=https%3A%2F %2Fapp.sp1.local%2F causes a System Error message, because the second query parameter separators are incorrect. The correct separator is an ampersand & and not a question mark?.



You must also use ampersands for all subsequent separators between additional query parameters in the URL.

In addition, you must URL-encode query parameter values that contain restricted characters. For information about URL encoding, see, for example, HTML URL-encoding Reference.

For both sample issues, update the URL of the IdP-initiated single sign-on (SSO) to the following URL.

https://sso.idp.local/idp/startSSO.ping?PartnerSpId=sp1&TargetResource=https%3A%2F %2Fapp.sp1.local%2F

Resolving service-related errors

Resolve the Unexpected System Error message and partnernot active status.

If a user encounters an Unexpected System Error message with a reference code, ask the user for the reference code and search for the value in the server log. The log message should help determine the root cause, which usually requires a configuration change.

If a user encounters a partner not active status, select Active in the Connection Status section and click **Save** on the **Activation & Summary** window for the connection.

Example

Example

Unexpected System Error

When a PingFederate identity provider (IdP) server receives a SAML AuthnRequest message through the redirect binding, but such SAML profile is not selected in the applicable service provider (SP) connection, PingFederate replies with an Unexpected System Error response with a reference code and logs an error message similar to the following entry.

```
2015-12-03 15:43:52,936 ERROR [org.sourceid.servlet.ErrorServlet]
 Top level error (ref#kwlqbn): javax.servlet.ServletException:
 org.sourceid.saml20.bindings.BindingException: Incoming binding
 urn:oasis:names:tc:SAML:2.0:bindings:HTTP-Redirect is not enabled for
 (SP) ::: sp1
```



In this sample log message, kwlqbn is the reference code.

Solution

Update the applicable SP connection to allow the Redirect binding for inbound messages from the SP. This works if the redirect binding is one of the mutually-agreed SAML bindings that both parties use. Alternatively, the SP can send SAML AuthnRequest messages through an allowable SAML binding based on the configuration of your SP connection.

Troubleshooting authentication policy issues

Authentication policies, an optional configuration in PingFederate, help implement complex authentication requirements. Having a complex policy or multiple policies can result in unintended runtime behaviors. The org.sourceid.util.log.PolicyTreeLogger logger makes it easier to troubleshoot such issues.

About this task

Identify and resolve authentication policy issues.

Steps

1. Enable debug messages for the org.sourceid.util.log.PolicyTreeLogger class.



Tip:

If you have enabled debug messages for the org.sourceid logger, you have already enabled debug messages for the org.sourceid.util.log.PolicyTreeLogger class. Skip to step 3.

- a. Edit the <pf install>/pingfederate/server/default/conf/log4j2.xml file.
- b. Look for Limit categories inside the Loggers element, as shown in bold in the following example.

```
<Loggers>
   <!-- ========== -->
   <!-- Limit categories -->
   <!-- ======== -->
   <!--
   <Logger name="org.sourceid.util.log.PolicyTreeLogger" level="DEBUG" />
</Loggers>
```

c. Uncomment the org.sourceid.util.log.PolicyTreeLogger logger, as shown in the following example. Note that the <!-- and --> markers have been removed. Example:

```
<Loggers>
   <!-- ========== -->
   <!-- Limit categories -->
   <!-- ======== -->
   <Logger name="org.sourceid.util.log.PolicyTreeLogger" level="DEBUG" />
```

```
</Loggers>
```

This change is activated within a minute. No restart of PingFederate is required.

2. Save any changes made.

For a clustered PingFederate environment, repeat these steps on each applicable node.

3. Repeat the request that demonstrates the authentication policy issue.



It is most useful to initiate this request in a browser without any cookies from prior sessions.

4. After you have replicated the issue, correlate server log messages using the PF cookie and tracking ID values. For more information, see *Troubleshooting runtime errors*.

Look for DEBUG messages from the org.sourceid.util.log.PolicyTreeLogger class. Example:

For example, suppose the tracking ID value is wXzQbS8MfHG40wpsQPiREIenJjc for a given request. The following server log messages demonstrate the authentication flow.

```
DEBUG [org.sourceid.util.log.PolicyTreeLogger] Policy 'General clients
 policy' | Selector | generalClients | Yes
DEBUG [org.sourceid.util.log.PolicyTreeLogger] Policy 'General clients
 policy' | Authn Source | idFirst
DEBUG [org.sourceid.util.log.PolicyTreeLogger] Policy 'General clients
 policy' | Authn Source | idFirst
DEBUG [org.sourceid.util.log.PolicyTreeLogger] Policy 'General clients
 policy' | Authn Source | idFirst | Rule | Alpha
DEBUG [org.sourceid.util.log.PolicyTreeLogger] Policy 'General clients
 policy' | Authn Source | idFirst | Alpha
DEBUG [org.sourceid.util.log.PolicyTreeLogger] Policy 'General clients
 policy' | Authn Source | https://sso.alpha.local:8031
DEBUG [org.sourceid.util.log.PolicyTreeLogger] Authn Policy Tree setting
 User ID from attribute 'subject' from Source type 'Adapter' and source
ID 'idFirst'
DEBUG [org.sourceid.util.log.PolicyTreeLogger] Policy 'General clients
 policy' | Authn Source | https://sso.alpha.local:8031 | Success
DEBUG [org.sourceid.util.log.PolicyTreeLogger] Policy 'General clients
 policy' | Authentication Policy Contract | APC | Finished
```



Note:

For readability, this sample ignores the time stamp and the tracking ID information. In other troubleshooting scenarios, such information can be valuable.

Log messages are interpreted as follows:

a. PingFederate finds an applicable policy named General clients policy. The first checkpoint is an OAuth Client Set Authentication Selector instance generalClients. PingFederate routes

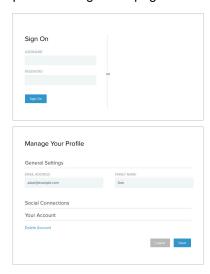
- this request to the Yes policy path because the client that submits the authorization request matches one of the clients defined authentication selector instance.
- b. PingFederate routes this request to an instance of the Identity First Adapter idFirst because that adapter instance is the next authentication source of the Yes policy path.
- c. Based on the user's provided user identifier, PingFederate determines that the Alpha rule applies and routes this request to the Alpha policy path.
- d. PingFederate routes this request to an identity provider (IdP) connection https://sso.alpha.local:8031 because that IdP connection is the next authentication source of the Alpha policy path. PingFederate also populates the subject attribute in the AuthnRequest message with the user identifier obtained from the Identity First Adapter instance.
- e. PingFederate receives a valid security token from the IdP https://sso.alpha.local:8031. PingFederate routes the request to the Success policy path, which ends with an authentication policy contract APC and concludes the authentication flow.

Troubleshooting registration and profile management issues

Certain browser extensions and their default settings can block a user's alternative authentication options configured to display on a third-party registration and profile management page. Review what browser a user currently employs and, if necessary, disable and fine-tune the extension settings of the browser.

About this task

Consider the following example. You have enabled third-party identity providers (IdP) to provide users with an alternative authentication option. Because you have enabled profile management, users can connect and disconnect their accounts with the third-party identity providers on the profile management page. However, some users are reporting that they do not see any such options on the sign-on page and the profile management page. The following screen captures illustrate what they see.



Browser extensions, particularly those designed to block ads or social network components, might block the user interface elements representing the third-party identity providers you enabled.

Steps

 Verify if users are using browser extensions that might have caused this issue. Disable or fine-tune the browser extension to resolve the issue.

Troubleshooting runtime errors

Users might encounter runtime errors when trying to connect to your partners. To troubleshoot these errors, investigate the user's activities, particularly the requests and the responses between the client and the PingFederate server, to determine the root cause.

Client side

Use built-in developer tools from the browsers to analyze HTTP traffic. Alternatively, you can use third-party tools, such as *Fiddler* and *Charles*.

Server side

Review server log messages to investigate what PingFederate has received from the client.

On rare occasions, you can also use third-party tools, such as Wireshark, or tools from the operating systems, such as *Tcpdump*, to capture network traffic on the PingFederate server.

For instructions of third-party tools, see their documentation.

To correlate server log messages to user activities, you can use one of the following:

- The tracking ID, which can be configured to show in the user-facing error pages in PingFederate.
- The PF cookie value, which requires capturing the HTTP headers on the client side.
- Real-time monitoring of the server log, which works well if the issues can be replicated reliably and involves using tools from the operating systems or third-party vendors.

Activating tracking ID in templates

You can configure PingFederate to display the tracking ID in the user-facing error Velocity templates. When an error occurs, use the tracking ID to look for the related log messages.

About this task

You can find the Velocity template files in the <pf install>/pingfederate/server/default/ conf/template directory.

The Velocity variable is \$TrackingId and is available in the following templates:

- general.error.page.template.html
- generic.error.msg.page.template.html
- idp.slo.error.page.template.html
- idp.sso.error.page.template.html
- sourceid-wsfed-idp-exception-template.html
- sp.slo.error.page.template.html
- sp.sso.error.page.template.html
- state.not.found.error.page.template.html

Steps

- 1. Open the applicable Velocity template file.
- 2. Search for the \$TrackingId variable.
- 3. Follow the inline instructions to activate the variable.



Note:

Template customization does not require a restart of PingFederate.

4. For a clustered PingFederate environment, repeat these steps on each engine node.

Result

The following screen capture demonstrates the user experience after the \$TrackingId variable is activated and an error has occurred. In this example, V3IwuUsy8PQp-9ZbE9UfUjOEo9c is the tracking ID.

Sign On Error

Unexpected Runtime Authn Adapter Integration Problem.

Please contact your system administrator for assistance regarding this error.

Adapter: IdpOpentoken

Tracking ID: tid:V3lwuUsy8PQp-9ZbE9UfUjOEo9c



Correlating log messages by PF cookie

If you do not want to activate the tracking ID in the user-facing error templates, you can capture the HTTP traffic and use the PF cookie value to find related server log messages for a given request.

About this task

Search the HTTP traffic server log and use the PF cookie value to review server log messages.

Steps

- 1. Capture HTTP traffic and look for the PF cookie value.
- 2. Search for the PF cookie value in the server log.
- 3. Because all server log messages, except the contents of the inbound requests and the outbound responses, are prefixed with their respective tracking IDs, use the tracking ID to review log messages and payloads pertaining to this transaction.

Result

In general, the most useful log messages are the ones tagged with WARN, ERROR, or prefixed with Caused by.

Example

Example

Suppose an error occurred and the associated the PF cookie value was

OaxBwPGw5OBeHVXe1sgifB7iZR5Rz2VI4rhJwqUSIXV. Based on the cookie value, you found the following log message.

2015-12-03 11:13:33,784 tid:V3IwuUsy8PQp-9ZbE9UfUjOEo9c DEBUG [org.sourceid.servlet.HttpServletRespProxy] adding lazy cookie Cookie{PF=OaxBwPGw5OBeHVXe1sgifB7iZR5Rz2VI4rhJwqUSIXV; path=/; maxAge=-1; domain=null} replacing null

After reviewing the related log messages based on the tracking ID V3IwuUsy8PQp-9ZbE9UfUjOEo9c, you found the next few messages.

2015-12-03 12:36:21,176 tid:V3IwuUsy8PQp-9ZbE9UfUjOEo9c ERROR [org.sourceid.saml20.profiles.idp.HandleAuthnRequest] Exception occurred

```
during request processing org.sourceid.websso.profiles.RequestProcessingException: Unexpected Runtime Authn Adapter Integration Problem.
```

. . .

Caused by: org.sourceid.saml20.adapter.AuthnAdapterException: Could not obtain attributes from the IdP Authentication Service.

Based on these log messages, the remedy is to review and update the configuration of the applicable identity provider (IdP) adapter instance.

Correlating log messages by tracking ID

All server log messages, except the contents of the inbound requests and the outbound responses, are prefixed with their respective tracking IDs, which helps with locating related log messages and payloads for a given transaction for troubleshooting.

About this task

Review the server log messages using the Tracking ID value provided by the user.

Steps

- 1. Ask the user for the Tracking ID value in the error message.
- 2. Search for the tracking ID in the server log.
- 3. Review the log messages and payloads pertaining to the transaction that is associated with the tracking ID.

Result

In general, the most useful log messages are the ones tagged with WARN, ERROR, or prefixed with Caused by.

Example

Example

Suppose an error occurred and the associated the tracking ID was V3IwuUsy8PQp-9ZbE9UfUjOEo9c. Based on the tracking ID, you found the following log message.

```
2015-12-03 11:13:33,784 tid:V3IwuUsy8PQp-9ZbE9UfUjOEo9c DEBUG [org.sourceid.servlet.HttpServletRespProxy] adding lazy cookie Cookie{PF=OaxBwPGw5OBeHVXe1sgifB7iZR5Rz2VI4rhJwqUSIXV; path=/; maxAge=-1; domain=null} replacing null
```

After reviewing the related log messages, you found the next few messages.

```
2015-12-03 12:36:21,176 tid:V3IwuUsy8PQp-9ZbE9UfUjOEo9c ERROR [org.sourceid.saml20.profiles.idp.HandleAuthnRequest] Exception occurred during request processing org.sourceid.websso.profiles.RequestProcessingException: Unexpected Runtime Authn Adapter Integration Problem.
```

Caused by: org.sourceid.saml20.adapter.AuthnAdapterException: Could not obtain attributes from the IdP Authentication Service.

Based on these log messages, the remedy is to review and update the configuration of the applicable identity provider (IdP) adapter instance.

When enabled on PingDirectory, you can correlate events in PingFederate with LDAP activities in PingDirectory by looking for the matching session and request tracking IDs in their logs.

PingFederate can receive many requests during a session. The session ID is consistent throughout a session, but the request ID is unique for every request. You can use the request ID to search for specific events within a session.

PingFederate records runtime requests in its audit log and gives them tracking IDs. When PingFederate sends an LDAP call to PingDirectory, PingFederate also sends the request's tracking ID.



Important:

For PingFederate to send the tracking ID to PingDirectory, the Intermediate Client Request Control (OID=1.3.6.1.4.1.30221.2.5.2) must be enabled in PingDirectory. Also, there cannot be any access control instructions that prevent the PingFederate service account accessing PingDirectory from using this OID.

PingDirectory records the tracking ID as a session ID or request ID value in its access log. In the log, the ID is a property of a via element.

For example, if you see the following via elements in the PingDirectory access log, you can match them with PingFederate events by looking for session ID kkLivppizq1RvbaYBAuB1r9z-Y8' and request ID FhM15Lz0KwsQphYU1UVHS4xkC5s in the PingFederate audit log.

```
via="app='PingFederate' sessionID='tid:kkLivppizq1RvbaYBAuB1r9z-Y8'"
via="app='PingFederate' requestID='tid:FhM15Lz0KwsQphYUlUVHS4xkC5s'"
```



When a PingFederate endpoint receives a request, it records a request ID at the DEBUG level of the server log. When the Request Header for Correlation ID field in the General Settings window specifies a request header, if the request includes that header, and the header's value contains 1 to 50 alphanumeric characters and hyphens, then PingFederate uses that value for the request ID. Otherwise, PingFederate uses a unique value that it generates for the request ID. For more information, see Configuring general settings on page 857.

Troubleshooting OAuth transactions

This troubleshooting guide walks through an OAuth request by inspecting the parameters provided by the client in different stages of an OAuth transaction. It then compares the parameter values against the corresponding settings defined in the PingFederate administrative console.

Troubleshooting OAuth use cases involves reviewing the OAuth requests and various OAuth settings.

While the guide focuses on an OAuth authorization code use case, in which the end user authenticates through an identity provider adapter, it provides a general guidance for other OAuth use cases with or without OpenID Connect in terms of which part of the configuration comes into play and how a request might fail at which stage.

Reviewing an OAuth request and various OAuth settings

To troubleshoot OAuth requests, inspect the parameters provided by the client in different stages of an OAuth transaction, and then compare the parameter values against the corresponding settings defined in the PingFederate administrative console.

About this task

A typical OAuth request looks like the following with the parameters that are submitted by a client. These are what you track as you go through the configuration during a troubleshooting task. Your requests could look very different depending on the specifics of your authorization server, resource server, and clients.

?client id=client&response type=code&redirect uri=uri&scope=scope1 scope2

In this example request, the client is providing 4 parameters:

- client_id
- response_type
- redirect_uri
- scope



Note:

There are other optional parameters that can be included in the request but are at this time not of interest to you in troubleshooting a typical request.

Steps

1. Review OAuth request in the server log.

PingFederate logs requests and responses to the server log. The details vary based on the log level set in <pf install>/pingfederate/server/default/conf/log4j2.xml file. The following example shows a client making an Authorization Code grant type, as indicated by the response_type parameter of code.

```
2015-11-29 22:11:35,795 tid:aUBgyTMjPfuSFp5BYtsK6Cb2McM DEBUG
[org.sourceid.websso.servlet.ProtocolControllerServlet] ---REQUEST
(GET)/as/authorization.oauth2 from 127.0.0.1:
    ---PARAMETERS---
    scope:
        list_users
    response_type:
        code
    client_id:
        pa_web_session_9
```

- 2. Check if the client is valid.
 - a. Go to Applications # OAuth # Clients to open the Clients window.
 - b. Look for the client by its Client ID.

Result:

If the client is not found, PingFederate denies the request, returns a 400 error to the client, and logs an <code>Unknown or invalid client_id message</code> to the server log, similar to the following code.

```
2015-11-29 22:11:35,812 tid:aUBgyTMjPfuSFp5BYtsK6Cb2McM DEBUG
[org.sourceid.oauth20.handlers.HandleAuthorizationRequest]
Normal exception being handled during OAuth request processing:
org.sourceid.oauth20.handlers.AuthorizationRequestException: Unknown or
invalid client_id
```

- 3. Check if the redirect_uri parameter value is valid for the client.
 - a. Go to Applications # OAuth # Clients to open the Clients window.
 - b. Go to the Clients window.
 - Select the applicable client.
 - d. Compare the redirect_uri parameter value against the values defined in the Redirect URIs field.

Result:

If the request comes without a redirect_uri parameter and the Redirect URIs field contains multiple entries, PingFederate denies the request, returns a 400 error to the client, and logs an Invalid redirect_uri message to the server log, similar to the following code.

```
2015-11-29 22:23:59,858 tid:aUBgyTMjPfuSFp5BYtsK6Cb2McM DEBUG [org.sourceid.oauth20.handlers.HandleAuthorizationRequest]
Normal exception being handled during OAuth request processing:
org.sourceid.oauth20.handlers.AuthorizationRequestException: Invalid redirect_uri
```

If the request comes with a redirect_uri parameter value that does not match any Redirect URIs values defined for the client, PingFederate denies the request, returns a 400 error to the client, and logs an Invalid redirect uri message to the server log.

- 4. Check if the **response type** parameter value is authorized for the client.
 - a. Go to Applications # OAuth # Clients to open the Clients window.
 - b. Select the applicable client.
 - c. In the Allowed Grant Types field, verify the response type is selected.

Result:

If the response_type value is not one of the allowable grant types, PingFederate denies the request, returns a 403 error to the client, and logs an unauthorized_client message with an error description to the server log, similar to the following code.

```
2015-11-29 22:25:51,212 tid:aUBgyTMjPfuSFp5BYtsK6Cb2McM DEBUG
[org.sourceid.saml20.bindings.LoggingInterceptor] Transported Response.
OutMessageContext:
    OutMessageContext
    entityId: pa_web_session_1 (null)
    virtualServerId: null
    Binding: oauth:authz
    params: {error=unauthorized_client, error_description=implicit grant
not allowed for this client}
    Endpoint: https://servapp.ext.den-ping.com/pa/oidc/
cb#error_description=implicit+grant+not+allowed+for+this
+client&error=unauthorized_client
    SignaturePolicy: BINDING_DEFAULT
```

- 5. Check if the scopes requested with the scope parameter are defined for the authorization server.
 - a. Go to System # OAuth Settings # Authorization Server Settings.
 - b. In the **Authorization Server Settings** window, compare the scopes requested against the values defined in the **Scope Value** or the **Scope Group Value** fields.

Result:

If the scopes requested are not defined, PingFederate denies the request, returns a 403 error to the client, and logs an <code>invalid_scope</code> message with an error description to the server log, similar to the following code.

```
2015-11-29 22:24:52,588 tid:aUBgyTMjPfuSFp5BYtsK6Cb2McM DEBUG
[org.sourceid.saml20.bindings.LoggingInterceptor] Transported Response.
OutMessageContext:
   OutMessageContext
   entityId: pa_web_session_1 (null)
   virtualServerId: null
   Binding: oauth:authz
   params: {error=invalid_scope, error_description=The requested
   scope(s) must be blank or a subset of the provided scopes.}
   Endpoint: https://servapp.ext.den-ping.com/pa/oidc/cb?
error_description=The+requested+scope%28s%29+must+be+blank+or+a+subset+of
+the+provided+scopes.&error=invalid_scope#.
   SignaturePolicy: BINDING_DEFAULT
```

- 6. Check if the scopes requested are valid for the client.
 - a. Go to Applications # OAuth # Clients to open the Clients window.
 - b. Select the applicable client.
 - c. If the client is limited to specific scopes, as indicated by the selection of the **Restrict Scopes** check box, verify the scopes requested are valid for the client.

Result:

If the scopes requested are not valid for the client, PingFederate denies the request, returns a 403 error to the client, and logs an invalid scope message with an error description to the server log.

7. Review the authentication process.

Suppose this OAuth request uses an identity provider (IdP) adapter for authentication. Check the IdP adapter mapping and the runtime selection made by the user.

- a. Go to Authentication # OAuth # IdP Adapter Grant Mapping.
- b. Verify an entry exists for the IdP adapter involved.

Result:

If more than one option is available, authentication policies can be used to select an authentication source. If no authentication policy is defined or applicable, the user is prompted with a list of all available authentication sources. The user can also save the preferred authentication source for later in the form of a pfidpaid cookie.

If a selection was made and the authentication source is not defined for OAuth, an error is returned to the user.

- 8. Upon successful authentication, PingFederate presents to the user an authorization consent page or a redirection to a trusted web application that is responsible to prompt the user for authorization unless a bypass option is configured. Review the authorization approval settings.
 - a. Go to System # OAuth Settings # Authorization Server Settings.
 - b. Review the Reuse Existing Persistent Access Grants for Grant Types setting.

Result:

If the grant type is selected, the authorization consent page is bypassed for the same client, the same user and same, or lesser, scope.

c. Review the Consent User Interface setting.

If PingFederate is configured to use an external consent user interface, verify that the associated settings are correctly configured and the web application is fulfilling its responsibilities.

- d. Go to Applications # OAuth # Clients to open the Clients window.
- e. Go to the Clients window.
- f. Select the applicable client.
- g. Review the Bypass Authorization Approval setting.

Result:

If the **Bypass Authorization Approval** check box is selected, the authorization consent page is bypassed as well.

9. Verify a mapping is configured.

When authorization is obtained, PingFederate maps attribute values from the authentication source into the persistent grants, the user_key, user_name, and extended attributes defined in the Authorization Server Settings window. This is the first stage of the two-stage OAuth attribute mapping process. In this example, because the user authenticates through an IdP adapter, check the IdP adapter mapping.

- a. Go to Authentication # OAuth # IdP Adapter Grant Mapping.
- b. Verify an entry exists for the IdP adapter involved and review its configuration.
- 10. Review the request and the settings related to access token management.

Finally, PingFederate selects the applicable access token management (ATM) instance and fulfill the access token by mapping values from the persistent grants, the authentication source, or both. This is the second stage of the two-stage OAuth attribute mapping process.

At runtime, the PingFederate OAuth authorization server uses the following rules to determine which ATM instances to use:

a. PingFederate limits the eligible ATM instances to those that are available in the context of the request. For most requests, these are instances that have an attribute mapping defined in the Access Token Mapping window. For OAuth Assertion Grant requests, it is the set of instances

- for which a mapping is defined in the IdP connection. If configured, the ACL can also limit which ATM instances are eligible.
- b. If the request comes with an access_token_manager_id or aud parameter, PingFederate uses the information to determine the applicable ATM instance.
- c. If the request does not come with either parameter, for OAuth clients supporting the OpenID Connect protocol by including the openid scope value, PingFederate uses the ATM instance specified by the OpenID Connect policy associated with the client. For resource server clients, you can optionally configure PingFederate to use any eligible ATM instances for the purpose of token validation.
- d. If the request comes with neither of the two parameters nor the <code>openid</code> scope, PingFederate uses the default ATM instance of the client if configured, or the default ATM instance defined for the installation if eligible. For token validation requests, if resource server clients do not provide either the <code>access_token_manager_id</code> or <code>aud</code> parameter in their requests and the resource server clients have not been configured to validate against any eligible ATM instances, the same logic applies.
- a. Determine if the OAuth request is sent to the /as/authorization.oauth2 authorization endpoint or the /as/token.oauth2 token endpoint with an access_token_manager_id or aud parameter.
- b. Go to Applications # OAuth # Clients to open the Clients window.
- c. Select the applicable client.
- Verify if a default access token is selected from the Default Access Token Manager list.
- e. Go to Applications # OAuth # Access Token Mapping.
- f. Review the attribute mapping configuration for the authentication source, if such mapping exists, or the Default mapping.

Other runtime issues

Possible runtime issues and their solutions related to unexpected certificate expiration and CrossModule and Network error messages.

Runtime issues and solutions

Problem	Solution
Certificates unexpectedly expire.	Verify that the server clocks are synchronized on both sides of the federation. You can configure PingFederate to notify administrators in advance of impending certificate expiration. For more information, see <i>Runtime notifications</i> on page 863.
Receive CrossModule or Network error messages when PingFederate is deployed with a supported hardware security module (HSM).	Verify network connections to the HSMs are active and running. Also ensure the HSMs have not been unintentionally shut down.

Collecting support data

When Ping Identity Support is helping an administrator troubleshoot a problem, it might ask the administrator to use the collect support data tool to compile information about your PingFederate installation.

About this task

The tool collects the following information by default:

- pingfederate/bin
- pingfederate/log (the most recent files of each type within a size limit)
- pingfederate/server/conf(configuration files)
- pingfederate/server/data (not key files)

The tool collects the following environment details:

- files present and their sizes
- certificate data
- version data
- JVM details
- and so on

The tool also collects the following system details, depending on the operating system:

- Crontab
- Ifconfig
- Netstat
- Uname
- and so on

If Ping Identity Support needs more information about the PingFederate installation than the default configuration provides, Support might ask the administrator to add a data collector to the tool by modifying its csd configuration.yaml file.

The tool consists of the following files in the pingfederate directory:

- bin/collect-support-data.bat
- bin/collect-support-data.sh
- bin/csd-1.0.jar
- server/default/conf/collect-support-data/csd configuration.yaml



You can use other parameters with the collect-support-data tool, including:

--encrypt

Encrypts the CSD output.

--outputPath

Specifies a directory for the CSD package.

For additional information about these parameters and more, use the --help command.

Steps

- 1. Using your PingFederate administrator account and a terminal, navigate to the pingfederate/bin directory.
- 2. Use one of the following commands to run the collect support data tool, depending on your operating system:

Choose from:

- On a Windows operating system, use collect-support-data.bat.
- On a Unix-based operating system, use ./collect-support-data.sh.

Result: As the tool collects data, it displays its progress and any errors. When it finishes collecting data, the tool places the data in a zip file in the bin directory. The file name format is supportdata-ping-\${hostname}-r-\${timestamp}.zip.

- 3. Review any errors that the tool displayed during the process or added to the log file support-dataping-\${hostname}-r.log. If needed, resolve the errors and run the tool again.
- 4. Send the support data zip file to Ping Identity Support according to the instructions from Support.

WS-Trust STS configuration

The PingFederate WS-Trust Security Token Service (STS) provides security-token validation and creation to extend single sign-on access to identity-enabled web services.

The section provides instructions for configuring the WS-Trust STS. For more information, see Web services standards on page 47.

Server settings

To use the PingFederate WS-Trust security token service (STS) for partner connections, enable and configure the WS-Trust protocol in your system server settings.

After you enable the protocol, you must identify the STS server with a unique federation identifier for both SAML 2.0 and SAML 1.1 tokens, unless these IDs are already established for the corresponding browserbased single sign-on protocols.

In addition, also within these server settings tasks, you have the option to require authentication globally for access to STS endpoints.

Enabling the WS-Trust protocol

To use the PingFederate WS-Trust security token service (STS) for partner connections, enable the WS-Trust protocol.

About this task

Using the administrative console, configure your server settings to enable the WS-Trust protocol.

Steps

- 1. Go to System # Server # Protocol Settings.
- 2. On the **Federation Info** tab, enter your SAML federation IDs unless these IDs are already established for the corresponding browser-based SSO protocols.



Note:

Identifiers are required for both SAML 2.0 and SAML 1.x to enable the STS to issue either type of token when requested. If you have not established a federation ID for either of these protocols or do not expect to use one or the other, enter a placeholder in any format and reconfigure later.

Configuring STS authentication

Configure PingFederate to require that client applications provide credentials to access the STS.

About this task

Although it is an optional configuration, configuring security token service (STS) authentication is recommended for identity provider (IdP) configurations that use the Username Token Processor. For other token processors and token generators, trust in the identity of the client is conveyed within the token itself and verified as part of processing. You can still configure authentication requirements to add another layer of security by limiting access to only authenticated clients.



Note:

You can configure STS authentication to either apply globally to all token formats and for all IdP and service provider (SP) partner connections, or token-to-token mappings, using more fine grained controls, at the connection level through issuance criteria.

Steps

- 1. Go to System # Server # Protocol Settings.
- 2. On the WS-Trust STS Settings tab, click Configure WS-Trust STS Authentication to open the WS-Trust STS Settings window.
- 3. On the Authentication Methods tab, select the Require HTTP Basic Authentication check box, the Require Mutual SSL/TLS Authentication check box, or both.

If both the Require HTTP Basic Authentication check box and the Require Mutual SSL/TLS Authentication check box are selected, all clients must provide credentials for both mechanisms.



Important:

If you select the Require Mutual SSL/TLS Authentication check box, you must configure a secondary PingFederate HTTPS port pf.secondary.https.port in the run.properties file. For more information, see *Configuring PingFederate properties* on page 752.

- 4. If you select the Require HTTP Basic Authentication check box, manage user accounts on the HTTP Basic Authentication tab.
 - a. Click Create User.
 - b. In the HTTP Basic Authentication, enter a user name in the username field and a password in the **password** field.. Repeat to create additional user accounts for your client applications.
 - c. Click Done.



Note:

On the HTTP Basic Authentication tab, you can also delete user accounts and update their passwords.

- 5. If you select the Require Mutual SSL/TLS Authentication check box, on the Mutual SSL Authentication tab, click Configure Mutual SSL Authentication.
 - a. On the Authentication Options tab, you can select the Restrict Access by Subject DN check box and the Restrict Access by Issuer Certificate check box. Click Next.
 - If both options are selected, the client certificate used for authentication to the STS endpoints must meet both sets of restrictions.
 - b. If you selected the **Restrict Access by Subject DN** check box, enter one or more subject DNs on the Allowed Subject DNs tab.



Note:

On the Allowed Subject DNs tab, you can edit or delete existing entries but you must keep at least one subject DN.

- c. Click Next.. When finished, click Save.
- d. If you selected the Restrict Access by Issuer Certificate check box, on the Allowed Issuer Certificates tab, from the Issuer Certificate list, select one or more client certificates.
- e. Click Add.

If you have not yet imported the client certificate, click Manage Certificates to do so.



On the Allowed Issuer Certificates tab, you can remove existing entries but you must keep at least one issuer.

- f. On the **Summary** tab, review your mutual SSL/TLS authentication settings. Click **Done**. Result: This will take you back to the WS-Trust STS Settings window.
- 6. When you finish configuring WS-Trust STS settings, on the Summary tab, review the configuration. To keep your changes, click Save.

Identity provider STS configuration

This section covers the identity provider (IdP) configuration for the PingFederate WS-Trust security token services (STS).

Managing token processors

The PingFederate Security Token Service (STS) uses token processors to validate incoming tokens and token requests.

About this task

You must configure at least one processor in order to set up an STS connection or a token-to-token mapping.

For more information about WS-Trust, see Web services standards on page 47.

PingFederate comes bundled with the following token processors:

- JWT Token Processor
- Kerberos Token Processor
- OAuth Bear Token Processor
- SAML 1.1 Token Processor
- SAML 2.0 Token Processor
- Username Token Processor

You can deploy additional token translators from *Ping Identity website*.

- 1. Go to Authentication # Token Exchange # Token Processors.
- 2. In the Token Processors window, choose from the following options.

Option	Description
Create New Instance	Configure a new instance
<existing instance="" link=""> under Instance Name</existing>	Modify an existing instance
Check Usage	Review the usage of an existing instance
Delete or Undelete	Remove an existing instance or cancel the removal



By default, PingFederate automatically checks multi-connection errors whenever you access this window. This verifies that configured connections are not adversely affected by changes made here.

If you experience noticeable delays in accessing this window, you can disable automatic connection validation. Go to System # Server # General Settings.

Selecting a token processor type

Begin creating a token-processor instance by choosing the processor type.

About this task

Steps

- 1. Go to AuthenticationToken ExchangeToken Processors.
- 2. To create a new instance, click Create New Instance to open a new Create Token Processor **Instance** configuration window and follow the following steps.



Note:

To select an existing token processor instance, select an instance in the **Instance Name** section.

- a. Click on the **Type** tab.
- b. Enter a name in the **Instance Name** field and ID in the **Instance ID** field.
- c. From the **Type** list, select the token-processor type.
- d. From the Parent Instance list, select a parent instance.

This is useful when you are creating an instance that is similar to an existing instance. The child instance inherits the configuration of its parent. In addition, you can override one or more settings during the rest of the setup. Select the Override ... check box and make the adjustments as needed in one or more subsequent windows.

3. Click Next.

Configuring a token processor instance

PingFederate configures multiple token processor instances. Select one that best fits your environment needs.

About this task

Depending on the selected token processor, the Instance Configuration tab presents you with different parameters.

- For configuring a specific token processor bundled with PingFederate, see one of the relevant following sections:
 - Configuring a Username Token Processor instance on page 955
 - Configuring a Kerberos Token Processor instance on page 955
 - Configuring an OAuth Token Processor instance on page 956
 - Configuring a JSON Web Token Processor instance on page 956
 - Configuring a SAML Token Processor instance on page 957
- For add-on processors, see the online documentation referenced in the download package.

Configuring a Username Token Processor instance

The integrated Username Token Processor accepts and validates username security tokens.

Steps

- 1. Go to Authentication # Token Exchange # Token Processors to open the Token Processors window.
- 2. Select on an existing token processor instance by clicking its name in the **Instance Name** section, or create a new instance by clicking **Create New Instance**.
 - Result: This will open the Create Token Processor Instance window configuration.
- 3. On the **Instance Configuration** tab, configure the basics of this token processor instances.
 - a. If you have not yet defined the desired Password Credential Validator instance, click **Manage**Password Credential Validators to do so.
 - b. Click **Add a new row to 'Credential Validators'** to select a credential-authentication mechanism instance for this adapter instance.
 - c. From the **Password Credential Validator Instance** list, select a Password Credential Validator instance. Click **Update**.

Add as many validators as necessary. Use the up and down arrows to adjust the order in which you want PingFederate to attempt credential authentication. If the first mechanism fails to validate the credentials, PingFederate moves sequentially through the list until credential validation succeeds. If none of the Password Credential Validator instances can authenticate the user's credentials, and the challenge retries maximum has been reached, the process fails.



Note:

If usernames overlap across multiple Password Credential Validator instances, this failover setup could lock out those accounts in their source locations.

d. In the Field Value section, enter a value in the Authentication Attempts.

When the number of login failures reaches this threshold, the user is locked out for a period of time.

The default value is 3.

Configuring a Kerberos Token Processor instance

The integrated Kerberos Token Processor accepts and validates Kerberos tokens through a configured Kerberos realm.

About this task

It supports authentication mechanism assurance from Active Directory (AD) domain service, making it possible to restrict access to users authenticating through specific mechanisms. For more information, see *Authentication mechanism assurance* on page 314.

Steps

- 1. Go to AuthenticationToken ExchangeToken Processors.
- On the Instance Configuration tab, select the applicable domain from the Domain/Realm Name list. An AD domain or a Kerberos realm must be configured for use with the Kerberos Token Processor. If the domain you want does not appear, click Manage Active Directory Domains/Kerberos Realms to add it. For more information, see Active Directory and Kerberos on page 915.



Note:

Kerberos tickets can be accepted from domains other than the domain configured in the token processor if there is a transient, two-way trust. This trust exists by default when domains are joined within a single server forest. For more information, see *Multiple-domain support* on page 916.

Configuring an OAuth Token Processor instance

The PingFederate STS provides validation for OAuth 2.0 bearer tokens. To use the OAuth Token Processor, you must first configure an Access Token Management (ATM) instance.

About this task

For more information about PingFederate OAuth authorization server and access token management, see About OAuth on page 75 and Access token management on page 549.

Steps

- 1. Go to Authentication # Token Exchange # Token Processors.
- 2. On the **Instance Configuration** tab, configure the basics of the token processor instance.
 - a. In the Access Token Manager row, from the Field Value list, select an ATM instance. If the desired ATM instance is not shown, click Manage Access Token Manager. Result:
 - The token processor instance uses the selected ATM instance to validate the OAuth bearer access tokens.
 - Optional: Select the Scope Value as Single String check box.

Result:

If selected, the scope value is returned as a single space-delimited set of string value. If it is not selected, scope values are returned as a multivalued attribute.

Configuring a JSON Web Token Processor instance

The PingFederate security token services (STS) provides validation for JSON web tokens (JWTs).

Steps

- 1. Go to Authentication # Token Exchange # Token Processors.
- 2. On the **Instance Configuration** tab, enter the required information. See the following table for information about each field.

JWT Token Processor instance field names and descriptions

Field	Description
JWKS Endpoint URI	The URI of the JWKS endpoint. A set of JSON Web Keys (JWK) are downloaded from this endpoint and used for JWT signature verification.
Issuer	A unique identifier for the issuer of the JWT.

Field	Description
Expiry Tolerance	The amount of time, in seconds, to allow for clock skew between servers. Valid range is 0 to 3600.

Configuring a SAML Token Processor instance

The integrated SAML (1.1 or 2.0) Token Processor accepts and validates SAML (1.1 or 2.0) security tokens. The PingFederate security token service (STS) validates digital signatures using all trusted certificate authorities (CAs) imported into PingFederate.

About this task

On the **Instance Configuration** tab, configure a SAML Token Processor instance.

You can restrict the signature verification process by subject distinguished names (DN), issuers, or both to limit the token requests accepted for this token processor instance.

You must indicate a unique identifier for the PingFederate STS. Token processor instances reject SAML tokens that do not contain the identifier in the audience element.

Steps

- Go to Authentication # Token Exchange # Token Processors.
- On the Instance Configuration tab, configure the basics of the token processor instance.
 - a. In the **Audience** row, in the **field value** field, enter the URI that uniquely identifies your federation gateway for this SAML protocol.
 - This is the federation ID for the STS for either SAML 1.1 or SAML 2.0 tokens, depending on which processor you are configuring.
 - b. Optional: Click Add a new row to 'Valid Certificate Issuer DNs' to enter one or more issuers.



Important:

If issuer DNs are specified here, then only those issuers are considered valid for verifying incoming digital signatures. Otherwise, all trusted certificate authorities (CAs) are used to verify signatures.

 c. Optional: Click Add a new row to 'Valid Certificate Subject DNs' to enter one or more subject DNs.



Important:

If subject DNs are specified here, then only those subject DNs are considered valid for verifying incoming digital signatures. Otherwise, all trusted certificate authorities (CAs) are used to verify signatures.

If you specify both issuer DNs and subject DNs, then the certificate used to validate signatures must match an entry in both lists.

If you provide no issuer DN and subject DN, then all certificates are treated as valid for purposes of verification.

Token processors allow administrators to add to a built-in list of user attributes that the processor returns from an incoming token.

About this task



The Extended Contract tab shows a different list of attributes under Core Contract, depending on the token processor selected.

Steps

- 1. Go to Authentication # Token Exchange # Token Processors.
- 2. On the Extended Contract tab, in the Extend the Contract field, enter the name of the desired attribute. Click Add.



Important:

For the OAuth Bearer Token Processor, added attributes must also be among those configured with the associated access token management instance.

Repeat these steps to add additional attributes.

Setting attribute masking

Mask attribute values that PingFederate logs from a processor instance at runtime.

Steps

- 1. Go to Authentication # Token Exchange # Token Processors.
- 2. On the Token Attributes tab, in the Mask Log Values section, select the check box for the attribute whose value you want to mask in logs.



Note:

If OGNL expressions might be used to map derived values into outgoing tokens and you want those values masked in logs as well, select the Mask All OGNL-expression Generated Log Values check box under the list of attributes.

Reviewing the token processor configuration

Review your token processor instance summary information.

About this task

On the **Summary** tab, review, amend, save, or discard the token processor configuration.

Steps

- To keep your changes, click Save.
- To amend your configuration, click the name of the corresponding tab and then follow the configuration wizard to complete the task.
- To discard your changes, click Cancel.

Managing STS request parameters

Configure PingFederate to act as a WS-Trust security token service (STS) by defining sets of Request for Security Token (RST) metadata parameters that can be used for mapping attributes into outgoing SAML security tokens.

About this task

After these request contracts are defined, you can make them available when configuring WS-Trust STS settings in the SP connections. For more information, see Selecting a request contract on page 964.

To manage your request contracts:

Steps

- Go to Authentication # Token Exchange # STS Request Parameters Choose from:
 - To configure a new set of request parameters, click Add New Request Contract.
 - To modify an existing request contract, select it by its name in the Contract Name column.
- To review the usage of an existing request contract, in the Action section, click Check Usage.
- To remove an existing request contract or to cancel the removal request, in the Action section, click Delete or Undelete.

Creating a request contract

Specify one or more parameters that will be included in request-security-tokens (RST) applicable to connection partners. You can make request contracts available for token-attribute mapping during partnerconnection configuration.

About this task

Identify the contract and define parameters that will be available in token requests associated with this contract for partner connections.

Steps

- 1. Go to Authentication # Token Exchange # STS Request Parameters.
 - Result: This will open the STS Request Parameters window configuration.
- 2. Click Add New Request Contract.

Result:

This will open the **Request Contract** configuration window.

3. Enter the contract name in the Contract Name field and the contract ID in the Contract ID field.



Note:

In the Request Contract configuration window, after you click Done, the Contract Name and Contract ID cannot be modified.

4. In the Parameter Name field, enter the parameter that will be included in the RSTs. Click Add. You must add at least one parameter. Repeat this step to add more parameters.



Note:

After the request is saved, you can add, modify, or remove parameters. You must keep at least one parameter.

5. Click Done.

Configuring SP connections for STS

Configure a security token service (STS) connection to a service provider (SP) partner either in conjunction with browser-based single sign-on (SSO) or independently.

Steps

- 1. Go to System # Server # Protocol Settings.
- 2. On the **WS-Trust STS** tab, configure an STS connection.

Configuring protocol settings for IdP STS

Specify the WS-Trust protocol details for web service clients related to this connection.

Steps

- 1. Go to Applications # Integration # SP Connections.
- 2. Click on an existing connection in the SP Connection column, or click Create Connection to configure a new SP connection to open the SP Connection configuration window.
- 3. On the WS-Trust STS tab, click Configure WS-Trust STS to open the WS-Trust STS configuration window.



Note:

The WS-Trust STS tab is only available after you enable the WS-Trust role on the Connection Type tab. For more information, see Configuring SP Connections for STS.

4. On the Protocol Settings tab, enter a URL for your partner's web service in the Partner Service Identifier field. Click Add.

This identifier compares to the AppliesTo element in the Requests for Security Token (RST) messages and can be either a complete URL or a base URL for matching variable ports or paths.

Repeat this step to add additional identifiers.

5. Select any of the following WS-Trust protocol setting options that are applicable to your use case.

Option	Description
OAuth Assertion Profiles	When selected, four additional token-type requests become available based on these OAuth grant types:
	 JWT Bearer Token grant type OAuth Access Token via JWT Bearer Token grant type SAML 2.0 Bearer Assertion grant type OAuth Access Token via SAML 2.0 Bearer Assertion grant type
	See STS OAuth integration on page 74 for more information on the use of these token-type requests.
Default Token Type	The default token type when a web service client (WSC) does not specify in the token request which token type the STS should issue. The choices are:
	SAML 2.0SAML 1.1SAML 1.1 for Office 365

Option	Description
	The default token type does not need to match the protocol selected for the browser-based SSO, if enabled, and does not apply to OAuth assertion profiles because those RST messages must contain the requested token type.
Generate Key for SAML Holder of Key Subject Confirmation Method	When selected, the STS generates a symmetric key to be used in conjunction with the "Holder of Key" (HoK) designation for the assertion's Subject Confirmation Method.
	For information about HoK assertions, see Web Services Security SAML Token Profile.
	This option does not apply to OAuth assertion profiles.
Encrypt SAML 2.0 Assertion	When selected, the STS encrypts the SAML 2.0 assertion. Applicable only to SAML 2.0 security token.
	This option does not apply to OAuth assertion profiles.

- 6. On the **Protocol Settings** tab, customize SAML messages and assertions for WS-Trust connections. Message customizations are OGNL expressions that allow you to customize the security token sent from PingFederate to the service provider (SP).
 - a. Click Show Advanced Customizations.
 - b. From the Message Type list, select a type option and enter an expression. The message type is used to override the message type returned from the OGNL expression.

The following tables describe the relationship between message type and available variables, and the corresponding class or interface information in Java.

SP connections SAML 2.0 message types and expressions

Message types	Available variables and classes/interfaces in Javadoc
AssertionType	#AssertionType
	org.sourceid.saml20.xmlbinding.assertion.AssertionType
	#AssertionTypes
	org.sourceid.saml20.xmlbinding.assertion.AssertionType[]
	#Attributes
	org.sourceid.util.log.AttributeMap

Message types	Available variables and classes/interfaces in Javadoc
ResponseDocument	#ResponseDocument
	For a connection with WS-Trust v1.3, #ResponseDocument will be of type org.oasisOpen.docs.wsSx.wsTrust.x200512. RequestSecurityTokenResponseCollectionDocument
	For a connection with WS-Trust v1.2, #ResponseDocument will be of type org.xmlsoap.schemas.ws.x2005.x02.trust. RequestSecurityTokenResponseDocument
	#Attributes
	org.sourceid.util.log.AttributeMap

SP Connections SAML 1.x message types and expressions

Message types	Available variables and classes/interfaces in Javadoc
AssertionType	#AssertionType
	org.sourceid.protocol.saml11.xml.AssertionType
	#AssertionTypes
	org.sourceid.protocol.saml11.xml.AssertionType[]
	#Attributes
	org.sourceid.util.log.AttributeMap
ResponseDocument	#ResponseDocument
	For a connection with WS-Trust v1.3, #ResponseDocument will be of type org.oasisOpen.docs.wsSx.wsTrust.x200512. RequestSecurityTokenResponseCollectionDocument
	For a connection with WS-Trust v1.2, #ResponseDocument will be of type org.xmlsoap.schemas.ws.x2005.x02.trust. RequestSecurityTokenResponseDocument
	#Attributes
	org.sourceid.util.log.AttributeMap

Example:

Example of an AssertionType expression for SAML1.1.

```
#AssertionType.getAuthenticationStatementArray(0)
.getSubject().getNameIdentifier().setStringValue("JoeSAML2IDP"),
#AssertionType
```

Example of a ResponseDocument expression for a connection with WS-Trust v1.3.

#RequestSecurityTokenResponseCollectionDocument.getRequestSecurityTokenResponseCo
.getRequestSecurityTokenResponseArray(0).setContext('context1'){code}

• Example of a **ResponseDocument** expression for a connection with WS-Trust v1.2.

#RequestSecurityTokenResponseDocument.getRequestSecurityTokenResponse().setContex

7. ClickNext.

Specify a token's timeframe of validity before and after issuance.

About this task

Standards require a window of time during which a security token is considered valid. Each token has a time-stamp XML element as well as elements indicating the allowable lifetime of the token, in minutes, before and after the token time stamp.

Steps

- 1. Go to Applications # Integration # SP Connections.
- 2. In the WS-Trust STS window, click the Token Lifetime tab.
- 3. Override the default values for the fields in the following table.

Token timeframe parameters and descriptions

Field	Description
Minutes Before	The amount of time before the SAML token was issued during which it is to be considered valid. The default value is 5.
Minutes After	The amount of time after the SAML token was issued during which it is to be considered valid. The default value is 30.

4. Click Next.

Configuring token creation

The PingFederate security token service (STS) requires creating tokens to enable web services access to resources at your service provider (SP) partner's site.

About this task

For the PingFederate STS to issue a security token in response to requests for partner services, you must indicate what user attributes are to be included in the token attribute contract. The attribute values sent in the token are then derived by mapping those available from the token processor you select. As with browser single sign-on (SSO), the mapping can be augmented using local data stores, variable or constant text, or expressions.

Steps

On the Token Creation, click Configure Token Creation to begin a token creation configuration.

Defining an attribute contract for IdP STS

During token creation configuration, define an attribute contract that the server sends in the security tokens issued in response to a web service client at your site.

About this task

An attribute contract is the set of user attributes that a web service client at your site expects to receive in security tokens issued for this connection. You identify these attributes on the **Attribute Contract** tab. For more information, see .

Steps

1. Enter the attribute name in the Extend the Contract field. Attribute names are case-sensitive and must correspond to the attribute names, including claims, expected by the requesting web services client (WSC).

Result:



The Format attribute associated with the NameID element in outgoing SAML tokens can be set by adding an attribute called **SAML NAME FORMAT**. The value of that attribute can then be mapped later. For more information, see Configuring contract fulfillment for token creation on page 969.

For information about the NameID elements and applicable URI values, locate the SAML 2.0 specification at www.oasis-open.org/standards.



You can add a special attribute, **SAML AUTHN CTX**, to indicate to the service provider (SP) the type of credentials used to authenticate to the identity provider (IdP) application-authentication context. Map a value for the authentication context on the attribute-mapping window later in the configuration, from any available attribute source, including the RST if a requested context is specified as a request parameter. For more information, see Configuring contract fulfillment for token creation on page 969.

2. Optional: For SAML 1.1 tokens, select a attribute namespace from the list.

This field appears only when the chosen default token type is SAML 1.1 or SAML 1.1 for Office 365 in the WS-Trust STS # Protocol Settings configuration.

Change the default namespace selection if you and your SP partner have agreed to a specific namespace.



Note:

You can customize name-format alternatives in the custom-name-formats.xml configuration file located in the <pf install>/pingfederate/server/default/data/config-store directory. You must restart PingFederate to activate any changes made to this file.

For more information about attribute namespace, see Attribute contracts on page 93.

- Click Add.
- 4. Repeat until all applicable attributes are defined.
- 5. Click Next.

Result

Use the Edit, Update, and Cancel workflow to make or undo a change to an item. Use the Delete and **Undelete** workflow to remove an item or cancel the removal request.

Selecting a request contract

Select a request contract to be used to map attribute values into the security token.

About this task

This optional setting allows you to use XML parameters contained in RSTs for token-attribute mapping. For more information, see *Managing STS request parameters* on page 959. If you are not using request parameters, click Next to continue.



Note:

If you are editing an existing configuration, you can change the request contract or disable this optional setting. These changes night require additional configuration changes in subsequent tasks.

Steps

- 1. To use request parameters, select the **Request Contract** check box to access the **Request** Contractlist.
- 2. Select a request contract from the Request Contract list.

If the contract you want is in the Request Contract list, click Manage STS Request Parameters.

When selected, you can choose the request contract as the attribute source in the IdP Token Processor Mapping configuration later in the setup.

3. Click Next.

Managing IdP token processor mappings

Identity provider (IdP) token processors are responsible for validating incoming security tokens as part of an STS operation. A configured and deployed token processor in PingFederate is known as a token processor instance. Map, edit, remove, or save a token processor instance.

About this task

You can map one or more token processor instances into an service provider (SP) connection to satisfy multiple session-management requirements where needed. The same token processor instances can be mapped in multiple SP connections.

When token processor instances are restricted to certain virtual server IDs, the allowed IDs are displayed in the Virtual Server IDs column.

Steps

- 1. In the Token Creation window, click the IdP Token Processor Mapping tab.
- 2. To map a token processor instance, click Map New Token Processor Instance. Choose from:
 - To edit the mapping configuration of a token process instance, open it by clicking on its name, select the setting that you want to reconfigure, and complete the change.
 - To remove a token processor instance or cancel the removal request, click **Delete** followed by Save or Undelete.
- 3. If you are creating a new connection and you are finished with mapping configuration, click **Done**.
- 4. If you are editing an existing configuration and want to keep your changes, click Save.

Selecting a token processor instance

Select an IdP token processor instance that can be used to authenticate users for a partner. Attributes returned by the token processor instance you choose for the token processor contract can be used to fulfill the attribute contract with your partner.

About this task

On the **Token Processor Instance** tab, choose an instance of a deployed token processor that suits your requirements for this connection.



Note:

Steps

1. From the **Token Processor Instance** list, select a token processor instance.



If you do not see the desired token processor instance, click Manage Token Processor Instances to create a new instance of any deployed token processor.

2. Select the Override Instance Settings check box if you want to customize one or more token processor settings for this connection alone.



Alternatively, you can create child token processor instances of a base token processor instance with overrides so that such customized settings can be applied to several connections. For more information, see Hierarchical plugin configurations on page 90.

Result:

When selected, the administrative console adds a new Override Instance tab.

Overriding a token processor instance

Customize the token processor instance settings for overriding one or more token processor settings for a connection.

About this task

On the Override Instance tab, configure override token processor settings specifically for this connection.



Note:

Any changes to the base token processor instance are propagated to a connection provided the same changes are not overridden for the connection.

Steps

- 1. Click Override Instance Settings.
- 2. On the Instance Configuration, Extended Contract, and Token Attributes tabs, for each of the settings, select the **Override** check box and make your changes.



Note:

If you are editing a currently mapped token processor instance, click Override Instance Settings to reconfigure any overridden settings for this connection. You can also remove all overridden settings on a per-tab basis by clearing the **Override** check box near the top of the window.

- Click Next.
- 4. Click **Done** to close the **Override Instance Settings** window and continue with the rest of the **IdP** Token Processor Mapping window configuration.

Virtual server IDs provide more configuration flexibility in cases where you need to identify your server differently when connecting to a partner in one connection for multiple environments or in multiple connections where the partner also supports multiple federation IDs. When you multiplex one connection for multiple environments, you can restrict each token processor added to a WS-Trust STS SP connection or IdP connection.

About this task

When you multiplex one connection for multiple environments see, you can enforce authentication requirements by restricting a token processor to certain virtual server IDs on the Virtual Server IDs tab. By default, no restriction is imposed.



If you are editing a currently mapped token processor instance, you can toggle the Restrict Virtual Server **IDs** setting. You can also change the allowed virtual server IDs.

Steps

- 1. In the IdP Token Processor Mapping configuration window, go to the Virtual Server IDs tab.
- 2. Select the Restrict Virtual Server IDs check box.
- 3. Select one or more virtual server IDs that you want to allow for this token processor.

Selecting an attribute retrieval method for token creation

For token creation, you can query local data stores to help fulfill the attribute contract in conjunction with attribute values supplied by the token processor you are using with PingFederate.

About this task

The values supplied by the token processor are shown on the **Attribute Retrieval** tab, in the **Token** Processor Contract section.

To determine whether you need to look up additional values, compare the attribute contract against the token processor contract or the request contract, if configured. If the attribute contract requires more information, determine whether local data stores can supply it.



Note:

If you are editing a currently mapped token processor instance, you can change the mapping method, which may require additional configuration changes in subsequent tasks.

Steps

- 1. Go to the IdP Token Processor Mapping window.
- 2. On the Attribute Retrieval tab, in the Token Processor Contract section, choose from the following options.

Choose from:

- Select the Retrieve Additional Attributes from Data Stores to Fulfill the Attribute Contract option if you want to configure one or more data stores to look up attribute for a single mapping.
- Select the Use only the Token Processor Contract Values in the Outgoing Token option if you do not require data store query.
- 3. When finished, click Next.

Configuring attribute sources and user lookup for token creation

About this task

Attribute sources are specific data store or directory locations containing information that might be needed for the attribute contract. They are used to retrieve supplemental attributes. You can use more than one attribute source when mapping values to the attribute contract. The order matters and affects the queries differently. For example, if you plan on using the result of a query as an input to a subsequent query, stack your attribute sources accordingly.



If you are editing a currently mapped token processor instance, you can add, remove, or reorder attribute sources, which might require additional configuration changes in subsequent tasks.

Steps

1. In the IdP Token Processor Mapping window, click the Attribute Sources & User Lookup tab.



The Attribute Sources & User Lookup tab is only visible if you selected the Retrieve Additional Attributes from Data Stores to Fulfill the Attribute Contract option on the Attribute Retrieval tab. For more information, see Selecting an attribute retrieval method for token creation.

2. Click Add Attribute Source.

Result: The Attribute Sources & User Lookup window configuration opens.

- 3. On the **Data Store** tab, choose a data store for PingFederate to look up attributes.
- 4. Enter a description in the Attribute Source Description field and a source ID in the Attribute Source **ID** field, if prompted, for the data store.
- 5. From the **Active Data Store** list, select a data store instance.



If the data store you want is not shown in the Active Data Store list, click Manage Data Stores to review or add a data store instance.

6. Depending on the data store type, the rest of the setup varies as follows.

Data store types and their required tasks

Data store type	Required tasks
JDBC	 Specifying database tables and columns on page 363 Entering a database search filter on page 364
LDAP	 Specifying directory properties and attributes on page 365 Defining encoding for binary attributes on page 367 (optional) Entering a directory search filter on page 367
Other	 Specifying data source filter and fields on page 369

- 7. Repeat steps 2 6 as needed.
- 8. Click Save to exit the Attribute Sources & User Lookup window configuration.

Map values to the attributes defined for the contract. These are the values that are included in the SAML security tokens sent to the service provider (SP).

Steps

In the Token Creation | IdP Token Processor Mapping window, on the Attribute Contract Fulfillment tab, for each attribute, select a source from the Source list and then choose or enter a value. You must map all target attributes.

Token

When selected, the **Value** list is populated with attributes from the token processor instance. Select the desired attribute from the list. At runtime, the attribute value from the token processor instance is mapped to the value of the attribute in the SAML security token.

For example, to map the value of the Username Token Processor's username attribute as the value of the TOKEN SUBJECT attribute on the contract, select Token from the Source list and username from the Value list.

Context

When selected, the Value list is populated with the available context of the transaction. Select the desired context from the list. At runtime, the context value is mapped to the value of the attribute in the SAML security token.



The HTTP Request and STS SSL Client Certificate Chain context values are retrieved as Java objects rather than text. For this reason, OGNL expressions are more appropriate to evaluate and return values, such as Expression.



Note:

When using the STS Basic Authentication Username, STS SSL Client Certificate's Subject DN, or STS SSL Client Certificate Chain contexts, ensure the associated authentication is enabled and configured on the System # Server # Protocol Settings # WS-Trust STS Settings tab.

Request

When selected, the Value list populates with parameter values from the token request received from the web service client. This selection is available only if a request contract was selected earlier on the Request Contract tab. Select the desired context from the list. At runtime, the context value is mapped to the value of the attribute in the SAML security token.

LDAP, JDBC, or Other

When selected, the Value list populates with attributes that you have selected in the Attribute Source & User Lookup window configuration. Select the desired attribute from the list. At runtime, the attribute value from the attribute source is mapped to the value of the attribute in the SAML security token.

Expression

When enabled, this option provides more complex mapping capabilities, such as transforming incoming values into different formats. Select Expression from the Source list, click Edit under Actions, and compose your OGNL expressions. All variables available for text entries are also available for expressions. For more information, see Text.

Expressions are not enabled by default. For more information about enabling and editing OGNL expressions, see .

No Mapping

Select this option to ignore the **Value** field.

Text

When selected, the text you enter is mapped to the value of the attribute in the single sign-on tokens at runtime. You can mix text with references to any of the values from the authentication source using the \${attribute} syntax.

You can also enter values from your data store, when applicable, using the following syntax.

```
${ds.attr-source-id.attribute}
```

where attr-source-id is the attribute source ID value and attribute is any of the selected attributes in the attribute source configuration.



If you are editing a currently mapped token processor instance, you can update the mapping configuration, which might require additional configuration changes in subsequent tasks.

Defining issuance criteria for token creation

About this task

On the **Issuance Criteria** tab, define the criteria that must be satisfied in order for PingFederate to process a request further. This token authorization feature provides the capability to conditionally approve or reject requests based on individual attributes.

Begin this optional configuration by adding a criterion. Choose the source that contains the attribute to be verified. Some sources, such as Mapped Attributes, are common to almost all use cases. Other sources, such as **JDBC**, depend on the type of configuration. Irrelevant sources are automatically hidden. After you select a source, choose the attribute to be verified. Depending on the selected source, the available attributes or properties vary. Finally, specify the comparison method and the desired, compared-to, value.

If you define multiple criteria, all criteria must be satisfied for PingFederate to move a request to the next phase. A criterion is satisfied when the runtime value of the selected attribute matches or does not match the specified value depending on the chosen comparison method. The multi-value contains and multi-value does not contain comparison methods are intended for attributes that might contain multiple values. Such criterion is considered satisfied if one of the multiple values matches or does not match the specified value. Values are compared verbatim. If you require complex evaluations, including conditional criteria or partial matching, define them using attribute mapping expressions.



Important:

When you multiplex one connection for multiple environments, consider using attribute mapping expressions to verify the virtual server ID in conjunction with other conditions, such as group membership information, to protect against unauthorized access. For more information, see Multiple virtual server IDs on page 106 and Issuance criteria and multiple virtual server IDs on page 219.



All criteria defined must be satisfied or evaluated as true for a request to move forward. As soon as one criterion fails, PingFederate rejects the request and returns an error message.

- 1. In the Token Generator Mapping & User Lookup configuration window, go to the Issuance Criteria
- 2. From the **Source** list, select the attribute's source.

Depending on the selection, the Attribute Name list populates with associated attributes. See the following table for more information.

Attributes or properties and descriptions

Source	Description
Context	Select to evaluate properties returned from the context of the transaction at runtime.
	Note:
	Because the HTTP Request and STS SSL Client Certificate Chain context values are retrieved as a Java object rather than text, use attribute mapping expressions to evaluate and return values.
JDBC, LDAP, or other types of datastore (if configured)	Select to evaluate attributes returned from a data source.
Mapped Attributes	Select to evaluate the mapped attributes.
Token	Select to evaluate attributes from the token processor instance.

3. Select the attribute to be evaluated under **Attribute Name**.



To evaluate the STS Basic Authentication Username, STS SSL Client Certificate Chain, or STS SSL Client Certificate's Subject DN context value, ensure that the associated authentication is enabled and configured on System # Server # Protocol Settings to open the WS-Trust STS Settings window.

4. From the **Condition** list, select the comparison method.

Available methods:

- equal to
- equal to (case insensitive)
- equal to DN
- not equal to
- not equal to (case insensitive)
- not equal to DN
- multi-value contains
- multi-value contains (case insensitive)
- multi-value contains DN
- multi-value does not contain
- multi-value does not contain (case insensitive)
- multi-value does not contain DN



Note:

The first six conditions are intended for single-value attributes. Use one of the multi-value ... conditions for PingFederate to validate whether one of the attribute values matches the specified value. When an attribute has multiple values, using a single-value condition causes the criteria to fail.



Note:

To evaluate the STS SSL Client Certificate's Subject DN context value, you must select one of the ... DN conditions. These methods normalize the DN before comparison to accommodate for different string representations that are still considered equivalent, such as case sensitivity or white space.

5. In the **Value** field, enter the comparison value.



Note:

Values are compared verbatim. If you require complex evaluations, including conditional criteria or partial matching, define them using attribute mapping expressions. For more information, see Attribute mapping expressions on page 214.

6. In the **Error Result** field, enter a custom error message.

The Error Result field is used by the faultstring element for SOAP 1.1 and the Reason/Text element for SOAP 1.2. For more information on SOAP, see Simple Object Access Protocol.

Using an error code in the Error Result field allows an application to process the code in a variety of ways, such as displaying an error message or e-mailing an administrator.

To use localized descriptions, enter a unique alias in the Error Result field, such as someIssuanceCriterionFailed. Insert the same alias with the desired localized text in the applicable language resource files, located in the <pf install>/pingfederate/server/ default/conf/language-packs directory.

If not defined, PingFederate returns ACCESS DENIED when the criterion fails at runtime.

- 7. Click Add.
- 8. Optional: Repeat to add more criteria.

- 9. If you require complex evaluations, including conditional criteria or partial matching, define them using attribute mapping expressions. For more information, see Attribute mapping expressions on page 214.
 - a. Click Show Advanced Criteria.
 - b. In the **Expression** field, enter the required expressions.
 - c. Optional: In the **Error Result** field, enter an error code or message.



Note:

If the expressions resolve to a string value instead of true or false, the returned value overrides the Error Result field value.

- d. Click Add.
- e. Optional: Click **Test**, enter values in the applicable fields, and verify the results.
- f. Optional: Repeat to add multiple criteria using attribute mapping expressions.

Reviewing the IdP token processor mapping

Review the identity provider (IdP) token processor mapping configuration to make and save changes as needed.

About this task

On the **Summary** tab, review, amend, save, or discard your configuration changes.

Steps

- To amend your configuration, click the corresponding tab title, then follow the configuration wizard to complete the task.
- To keep your changes, click **Done** and continue with the rest of the configuration.



When editing an existing configuration, you can also click Save as soon as the administrative console offers the opportunity to do so.

To discard your changes, click Cancel.

Selecting a request error handling method

If you are using request parameters to fulfill the attribute contract and the parameter values are not supplied in the request token (RST) messages, you can choose whether to continue or abort the tokencreation process.

About this task

On the **Error Handling** tab, select a request error handling method.



Note:

On the Request Contract tab, you must select the Request Contract check box and configure a request contract for this connection for the Error Handling tab to appear after the IdP Token Processor Mapping tab.

Steps

- 1. Go to **Applications** # **Integrations** # **SP Connections** to choose the appropriate connection.
- 2. On the WS-Trust STS tab, click Configure WS-Trust STS.
- 3. On the **Token Creation** tab, click **Configure Token Creation**.

- 4. On the **Error Handling** tab, select one of the following options. Choose from:
 - To allow the STS transaction to continue with null attribute values sent in the generated token, select Send User to SP Using Null Values for Attributes .
 - To abort the STS transaction with null attribute values sent in the generated token, select Abort the STS Transaction.

5. Click Next.

Reviewing the token creation configuration

Review the token creation configuration to make changes as needed.

About this task

On the **Summary** tab, review, amend, save, or discard your token creation configuration changes.

Steps

- To amend your configuration, click the corresponding tab title, then follow the configuration wizard to complete the task.
- To keep your changes, click **Done** and continue with the rest of the configuration.



Tip:

When editing an existing configuration, you can also click Save as soon as the administrative console offers the opportunity to do so.

To discard your changes, click Cancel.

Reviewing the IdP STS settings

Review the identity provider (IdP) security token service (STS) settings to make and save changes as needed.

About this task

On the **Summary** tab, review, amend, save, or discard the IdP STS settings.

Steps

- To amend your configuration, click the corresponding tab title, then follow the configuration wizard to complete the task.
- To keep your changes, click **Done** and continue with the rest of the configuration.



When editing an existing configuration, you can also click **Save** as soon as the administrative console offers the opportunity to do so.

To discard your changes, click Cancel.

Service provider STS configuration

This section covers the service provider (SP) configuration for security token service (STS).

Managing token generators

The PingFederate security token service (STS) uses token generators to issue security tokens that can be consumed by web services at your site. You must configure at least one generator in order to set up an STS connection or a token-to-token mapping.

About this task

In the **Token Generators** window, create, modify, review, or remove a token generator instance.



Note:

PingFederate comes bundled with the SAML 1.1 Token Generator and SAML 2.0 Token Generator.

You can deploy additional token translators from *Ping Identity website*.

For simplicity, this topic focuses on configuring an instance of the SAML 1.1 or 2.0 Token Generator. For add-on token generators, please refer to the online documentation referenced in the download package.

For more information about WS-Trust, see Web services standards on page 47.

Steps

- 1. Go to Applications # Token Exchange # Token Generators.
- 2. In the **Token Generators**window, choose from the following options.

Option	Description
Create New Instance	Configure a new instance
<existing instance="" link=""> under Instance Name</existing>	Modify an existing instance
Check Usage	Review the usage of an existing instance
Delete or Undelete	Remove an existing instance or cancel the removal



By default, PingFederate automatically checks multi-connection errors whenever you access this window. This verifies that configured connections are not adversely affected by changes made here.

If you experience noticeable delays in accessing this window, you can disable automatic connection validation. Go to System # Server # General Settings.

Selecting a token generator type

Token generators issue tokens for security token service and OAuth token exchange use cases. Select instances of a token generator type to use in mapping and policy configuration.

About this task

The first step in creating a token-generator instance is choosing the generator type.

Steps

- 1. Go to Applications # Token Exchange # Token Generators.
- 2. Select an existing generator instance by clicking its name or click Create New Instance to open the Create Token Generator Instance window.

- 3. On the **Type** tab, enter a name in the **Instance Name** field and ID in the **Instance ID** field for the token generator instance.
- 4. From the **Type**list, select the token-generator type.
- 5. From the **Parent Instance** list, select an option.

This is useful when you are creating an instance that is similar to an existing instance. The child instance inherits the configuration of its parent. You have the option to override one or more settings during the rest of the setup. Select the **Override** ... check box and make the adjustments as needed in one or more subsequent tabs.

6. Click Next.

Configuring a token generator instance

Configure the SAML token generator instance parameters depending on the use case requirements.

About this task

Depending on the selected token generator, the **Instance Configuration** tab presents you with different parameters.

Steps

- 1. Go to Applications # Token Exchange # Token Generators.
- 2. Select an existing generator instance by clicking its name or click **Create New Instance** to open the **Create Token Generator Instance** window.
- 3. On the **Instance Configuration** tab, configure the parameters for this instance type. For the integrated SAML 1.0 and 2.0 Token Generators, see the following table and specify parameters for generated SAML tokens.

SAML token generator instance field names and descriptions

Field	Instructions
Minutes Before	Enter a numerical value. This element in a SAML token allows for any server clock variability.
Minutes After	Enter a numerical value. This element in a SAML token allows for any server clock variability.
Issuer	Enter your SAML 2.0 entity ID or the SAML 1.x issuer as configured in the System # Server # Protocol Settings window.
Signing Certificate	Responses containing SAML tokens must be signed. Select a signing certificate from the list.
	If you have not yet created or imported your certificate into PingFederate, click Manage Signing Certificates . For more information, see <i>Manage digital signing certificates and decryption keys</i> on page 608.
Signing Algorithm	Select the signing algorithm corresponding to the selected certificate. Choices include SHA1 for both RSA and DSA, RSA-SHA256, SHA384, and SHA512, as well as, ECDSA-SHA256, SHA384, and SHA512.
Include Certificate in KeyInfo	If selected, the entire public certificate is included with the assertion. Otherwise, a short hash reference to the certificate is sent.
Include Raw Key in KeyValue	If selected, the raw key is included in the KeyInfo element as well.
Audience	A unique identifier for the target web service, used for the audience element of the generated SAML token.

Field	Instructions
Confirmation Method	Choose from among available methods:
	 urncm:sender-vouches Default option. urncm:bearer urncm:holder-of-key
	For more information, see WSS SAML Token Profile.
Encryption Certificate	The web service provider's public certificate for encryption is required only if holder-of-key is selected as the confirmation method. Select a partner certificate from the list.
	If you have not yet imported the certificate from your partner, click Manage Certificates to do so. For more information, see <i>Managing certificates from partners</i> on page 622.
Message Customization expression	Click Show Advanced Fields to see this field.
	An OGNL expression to customize the assertion. The returned type from the expression must be an AssertionType, or the customization will be ignored. The following example is for SAML2.
	<pre>#AssertionType.getSubject().getNameID().setStringValue(#AssertionType</pre>
	The following example is for SAML1.1.
	#AssertionType.getAuthenticationStatementArray(0).getSu #AssertionType

For information about add-on generators, see *Integrations Overview*.

4. Click Next.

Extending a token generator contract

Token generators allow administrators to add to a built-in list of user attributes that the generator includes in the outgoing token—an extended generator-attribute contract.

About this task

The **Extended Contract** tab shows a different list of attributes in the **Core Contract** section, depending on the token generator selected.

Steps

- 1. Go to Applications # Token Exchange # Token Generators.
- 2. Select an existing generator instance by clicking its name or click **Create New Instance** to open the **Create Token Generator Instance** window.
- 3. On the **Extended Contract** tab, in the **Extend the Contract** field, enter a name of the desired attribute. Click **Add**.
 - Repeat this step as needed to add another attribute.
- 4. Click Next.

Reviewing the token generator configuration

Review, save, or make changes as needed to your token generator instance configuration.

About this task

On the **Summary** tab, review, amend, save, or discard your token generator instance configuration.

Steps

- To keep your changes, click Save.
- To amend your configuration, click the name of the corresponding tab and then follow the configuration. wizard to complete the task.
- To discard your changes, click Cancel.

Configuring IdP connections for STS

Configure a security token services (STS) connection to an identity provider (IdP) partner either in conjunction with browser-based single sign-on (SSO) or independently.

Configuring protocol settings for SP STS

Configure the processing options for validating incoming SAML tokens in your identity provider (IdP) partner connection either in conjunction with browser-based single sign-on (SSO) or independently.

About this task

Select whether the STS should validate incoming tokens only or validate and then generate other types of tokens.

Steps

- 1. Go to Authentication # Integration # IdP Connections.
- 2. On the WS-Trust STS tab, click Configure WS-Trust STS.
- 3. On the Protocol Settings tab, from the Request Processing Options list, choose one of the following options:

Choose from:

- To only validate incoming SAML tokens, select Validate Incoming SAML Token.
- To validate and then also generate local tokens to enable single sign-on (SSO) access to web services at your site, select Validate Incoming SAML Token and Issue Local Token.



Note:

If you choose to generate local tokens as well, you must set up at least one token generator.

Configuring token generation

About this task

Details of this configuration are handled under the **Token Generation** task.

Steps

- Go to Authentication # Integration # IdP Connections.
- On the WS-Trust STS tab, click Configure WS-Trust STS.
- On the Token Generation tab, click Configure Token Generation.

Result: The **Token Generation** window opens.

Defining an attribute contract for SP STS

An attribute contract is the set of user attributes expected in incoming SAML assertions. For more information, see Attribute contracts on page 93.

About this task

On the Attribute Contract tab, identify the user attributes.

Optionally, you can mask the values of attributes, other than SAML SUBJECT, in logs that PingFederate writes when it receives security tokens.



Use the Edit, Update, and Cancel workflows to make or undo a change to an item. Use the Delete and **Undelete** workflows to remove an item or cancel the removal request.

Steps

- 1. Go to Authentication # Integration # IdP Connections.
- 2. On the WS-Trust STS tab, click Configure WS-Trust STS.
- 3. On the **Token Generation** tab, click **Configure Token Generation**.
 - Result: The **Token Generation** configuration window opens.
- Click the Attribute Contract tab.
- 5. Enter the name in the Extend the Contract field.



Note:

Attribute names are case-sensitive and must correspond to the attribute names expected by the requester.

- 6. Optional: Select the Mask Values in Log check box .
- 7. Click Add.
- 8. Repeat until all applicable attributes are defined.
- 9. Click Next.

Managing SP token generator mappings

Token generators provide a mechanism through which PingFederate can generate a local token based upon an incoming SAML token, including mapping user attributes to be included in the generated token. A configured and deployed token generator in PingFederate is known as a token generator instance.

About this task

In the Token Generator Mapping & User Lookup configuration window, manage service provider (SP) token generator mappings.

Map one or more token generator instances into an identity provider (IdP) connection to satisfy different token requirements by the web services at your site. The same token generator instances can be mapped in multiple connections.

When token generator instances are restricted to certain virtual server IDs, the allowed IDs are displayed under Virtual Server IDs.

Steps

1. Go to Authentication # Integration # IdP Connections # WS-Trust STS.

2. Click Configure WS-Trust STS.

Result: The WS-Trust STS configuration window opens.

3. On the Protocol Settings tab, from the Request Processing Options list, select Validate Incoming SAML Token and Local Issue Token. Click Next.

Result: This will add a **Token Generation** tab.

4. On the **Token Generation** tab, click **Configure Token Generation**.

Result: The **Token Generation** configuration window opens.

- 5. On the Token Generator Mapping & User Lookuptab, click Map New Token Generator Instance to open the Token Generator Mapping & User Lookup configuration window. Choose from:
 - To map a token generator instance, click Map New Token Generator Instance.
 - To edit the mapping configuration of a token generator instance, on the **Token Generator** Instance tab, click Manage Token Generator to open the Token Generators window. Click the token generator instance in the **Instance Name** section to open the configuration summary for this token generator instance. Select the setting that you want to reconfigure, and complete the change by clicking **Done**.
 - To remove a token generator instance or cancel the removal request, click **Delete** followed by Save or Undelete.
- 6. If you are creating a new connection and you are finished with mapping configuration, click Done. If you are editing an existing configuration and want to keep your changes, click Save.

Selecting a token generator instance

Select the token generator instance you would like to activate for incoming SAML messages from the service provider (SP) partner.

About this task



Note:

If you are editing a currently mapped token generator instance, you can toggle the **Override Instance Settings** setting. Clearing it removes all previously overridden settings for this connection. Selecting it provides you the opportunity to customize token processor settings specifically for this connection.

Steps

- 1. On the Token Generator Mapping & User Lookup tab, click Map New Token Generator Instance. Result: The **Token Generator Mapping & User Lookup** configuration window opens.
- 2. On the **Token Generator Instance** tab, select an instance from the **Token Generator Instance**.



Note:

If you do not see the desired token generator instance, click Manage Token Generator Instances to create a new instance of any deployed token generator.

Result:

When selected, the administrative console adds a **Override Instance** tab and a its new set of sub



You can also create child token processor instances of a base token processor instance with overrides so that the customized settings can be applied to several connections. For more information, see Hierarchical plugin configurations on page 90.

4. Click Next.

Overriding a token generator instance

Make changes to the token generator instance to override it or to leave it as is and propagate it for a particular service provider connection.

About this task

Override token generator settings for a specific service provider connection.



Note:

Any changes to the base token generator instance are propagated to a connection provided the same changes are not overridden for the connection.

Steps

- 1. On the Token Generator Mapping & User Lookup tab, click Map New Token Generator Instance. . Result: The **Token Generator Mapping & User Lookup** configuration window opens.
- 2. On the Override Instance tab, select the Override Instance Settings check box.



Note:

The override setting windows are functionally identical to those used for creating a new token generator instance. For more information, see *Managing token generators* on page 975.

- 3. On each of the settings windows, select the **Override** check box, make your changes, and then click Next.
- 4. When you are finished, click **Done** to continue with the rest of the mapping configuration.

Restricting a token generator to certain virtual server IDs

Virtual server IDs provide more configuration flexibility in cases where you need to identify your server differently when connecting to a partner in one connection for multiple environments or in multiple connections where the partner also supports multiple federation IDs.

About this task

When you multiplex one connection for multiple environments, you can enforce integration requirements by restricting a token generator to certain virtual server IDs on the Virtual Server IDs tab. By default, no restriction is imposed. For more information, see .



Note:

If you are editing a currently mapped token generator instance, you can toggle the Restrict Virtual Server **IDs** setting. You can also change the allowed virtual server IDs.

Steps

- 1. On the Token Generator Mapping & User Lookup tab, click Map New Token Generator Instance to open the Token Generator Mapping & User Lookup configuration window.
- 2. Click the Virtual Server IDs tab.
- 3. Select the Restrict Virtual Server IDs check box.
- 4. Select one or more virtual server IDs that you want to allow for this token generator.

Selecting an attribute retrieval method for token generation

You can fulfill the token generator contract by using only the attributes from the incoming SAML token or by using these attributes to look up additional information from a local data store.

About this task

For token generation, you can guery local data stores to help fulfill the token generator contract, in conjunction with attribute values supplied by the incoming token.

The values supplied by the token are shown in the Attribute Contract section on the Attribute Retrieval tab.

Steps

- 1. On the Token Generator Mapping & User Lookup tab, click Map New Token Generator Instance. Result: The **Token Generator Mapping & User Lookup** configuration window opens.
- 2. On the Attribute Retrieval tab, select how you want to fulfill the token generator contract for an instance.

Choose from:

- If the incoming SAML token contains all the attributes that your application requires, select Use only the attributes available in the incoming token.
- To set up a data store query, select Use the incoming token to look up additional information and then follow a series of sub tasks to complete the configuration.

For step-by-step instructions, see *Choosing a datastore* on page 363.



Note:

If you are editing a currently mapped token generator instance, you can change the mapping method, which might require additional configuration changes in subsequent tasks.

Configuring contract fulfillment for token generation

Fulfill your token generator contract requirements with values from the incoming SAML token, dynamic text, expressions, or from a data store lookup.

About this task

Map the values that the web services require to the attributes defined for the contract.

Steps

1. On the Token Generator Mapping & User Lookup tab, click Map New Token Generator Instance. Result: The **Token Generator Mapping & User Lookup** configuration window opens.

Assertion

When selected, the **Value** list populates with attributes from the incoming SAML token (assertion). Select the desired attribute from the list. At runtime, the attribute value from the assertion is mapped to the value of the attribute in the local token.

For example, to map the value of **TOKEN_SUBJECT** from a SAML assertion as the value of the **subject** user identifier on the token generator contract, select **Assertion** from the **Source** list and **TOKEN_SUBJECT** from the **Value** list.

Context

When selected, the **Value** list populates with the available context of the transaction. Select the desired context from the list. At runtime, the context value is mapped to the value of the attribute in the local token.



Note:

Because the HTTP Request and STS SSL Client Certificate Chain context values are retrieved as Java objects rather than text, use OGNL expressions to evaluate and return values, sees Expression.



Note:

When using the STS Basic Authentication Username, STS SSL Client Certificate's Subject DN, or STS SSL Client Certificate Chain contexts, ensure the associated authentication is enabled and configured on the System # Server # Protocol Settings # WS-Trust STS Settings tab.

• LDAP, JDBC, or Other

When selected, the **Value** list populates with attributes that you have selected from the data store. Select the desired attribute from the list. At runtime, the attribute value from the data store is mapped to the value of the attribute in the local token.

Expression

When enabled, this option provides more complex mapping capabilities, such as transforming incoming values into different formats. Select **Expression** from the **Source** list, click **Edit** under

Actions, and compose your OGNL expressions. All variables available for text entries are also available for expressions. For more information, see Text.

Expressions are not enabled by default. For more information about enabling and editing OGNL expressions, see .

No Mapping

Select this option to ignore the **Value** field.

Text

When selected, the text you enter is used at runtime. You can mix text with references to any of the values from the SAML token, using the \${attribute} syntax.

You can also enter values from your datastore, when applicable, using this syntax:

```
${ds.attribute}
```

where attribute is any of the attributes that you have selected from the data store.



Note:

If you are editing a currently mapped token generator instance, you can update the mapping configuration, which might require additional configuration changes in subsequent tasks.

3. Click Next.

Defining issuance criteria for token generation

PingFederate can evaluate various criteria to determine whether users are authorized to access service provider (SP) resources.

About this task

On the Issuance Criteria tab, define the criteria that must be satisfied in order for PingFederate to process a request further. This token authorization feature provides the capability to conditionally approve or reject requests based on individual attributes.

Begin this optional configuration by adding a criterion. Choose the source that contains the attribute to be verified. Some sources, such as Mapped Attributes, are common to almost all use cases. Other sources, such as JDBC, depend on the type of configuration. Irrelevant sources are automatically hidden. After you select a source, choose the attribute to be verified. Depending on the selected source, the available attributes or properties vary. Finally, specify the comparison method and the desired, compared-to, value.

If you define multiple criteria, all criteria must be satisfied for PingFederate to move a request to the next phase. A criterion is satisfied when the runtime value of the selected attribute matches or does not match the specified value depending on the chosen comparison method. The multi-value contains and multi-value does not contain comparison methods are intended for attributes that might contain multiple values. Such criterion is considered satisfied if one of the multiple values matches or does not match the specified value. Values are compared verbatim. If you require complex evaluations, including conditional criteria or partial matching, define them using attribute mapping expressions.



Important:

When you multiplex one connection for multiple environments, consider using attribute mapping expressions to verify the virtual server ID in conjunction with other conditions, such as group membership information, to protect against unauthorized access. For more information, see Multiple virtual server IDs on page 106 and Issuance criteria and multiple virtual server IDs on page 219.



Note:

Steps

- 1. On the Token Generator Mapping & User Lookup tab, click Map New Token Generator Instance. Result: The **Token Generator Mapping & User Lookup** configuration window opens.
- 2. Click the Issuance Criteria tab.
- 3. From the **Source** list, select the attribute's source.

Depending on the selection, the Attribute Name list populates with associated attributes. See the following table for more information.

Source entries and descriptions

Source	Description
Context	Select to evaluate properties returned from the context of the transaction at runtime.
	Note:
	Because the HTTP Request and STS SSL Client Certificate Chain context values are retrieved as a Java object rather than text, use attribute mapping expressions to evaluate and return values.
JDBC, LDAP, or other types of datastore (if configured)	Select to evaluate attributes returned from a data source.
Mapped Attributes	Select to evaluate the mapped attributes.
Token	Select to evaluate attributes from the incoming SAML token.

4. Select the attribute to be evaluated in the **Attribute Name** column.



To evaluate the STS Basic Authentication Username, STS SSL Client Certificate Chain, or STS SSL Client Certificate's Subject DN context value, ensure that the associated authentication is enabled and configured on System # Server # Protocol Settings to open the WS-Trust STS Settings window.

Available methods:

- equal to
- equal to (case insensitive)
- equal to DN
- not equal to
- not equal to (case insensitive)
- not equal to DN
- multi-value contains
- multi-value contains (case insensitive)
- multi-value contains DN
- multi-value does not contain
- multi-value does not contain (case insensitive)
- multi-value does not contain DN



Note:

The first six conditions are intended for single-value attributes. Use one of the multi-value ... conditions for PingFederate to validate whether one of the attribute values matches the specified value. When an attribute has multiple values, using a single-value condition causes the criteria to fail.



Note:

To evaluate the STS SSL Client Certificate's Subject DN context value, you must select one of the ... DN conditions. These methods normalize the DN before comparison to accommodate for different string representations that are still considered equivalent, such as case sensitivity or white space.

6. In the **Value** field, enter the comparison value.



Note:

Values are compared verbatim. If you require complex evaluations, including conditional criteria or partial matching, define them using attribute mapping expressions. For more information, see Attribute mapping expressions on page 214.

7. In the **Error Result** field, enter a custom error message.

The Error Result field is used by the faultstring element for SOAP 1.1 and the Reason/Text element for SOAP 1.2. For more information on SOAP, see the World Wide Web Consortium's Simple Object Access Protocol.

Using an error code in the Error Result field allows an application to process the code in a variety of ways, such as display an error message or e-mail an administrator.

To use localized descriptions, enter a unique alias in the Error Result field, such as someIssuanceCriterionFailed. Insert the same alias with the desired localized text in the applicable language resource files, located in the <pf install>/pingfederate/server/ default/conf/language-packs directory.

If not defined, PingFederate returns ACCESS DENIED when the criterion fails at runtime.

- 8. Click Add.
- 9. Optional: Repeat to add more criteria.

- 10. If you require complex evaluations, including conditional criteria or partial matching, define them using attribute mapping expressions. For more information, see Attribute mapping expressions on page 214.
 - a. Click Show Advanced Criteria.
 - b. In the **Expression** field, enter the required expressions.
 - c. Optional: In the **Error Result** field, enter an error code or message.



Note:

If the expressions resolve to a string value instead of true or false, the returned value overrides the Error Result field value.

- d. Click Add.
- e. Optional: Click **Test**, enter values in the applicable fields, and verify the results.
- f. Optional: Repeat to add multiple criteria using attribute mapping expressions.
- 11. When finished, click Next.

Reviewing the SP token generator mapping

Review the service provider (SP) token generator mapping instance.

About this task

On the **Summary** tab, review, amend, save, or discard your changes.

Steps

- To amend your configuration, click the corresponding tab title, then follow the configuration wizard to complete the task.
- To keep your changes, click **Done** and continue with the rest of the configuration.



When editing an existing configuration, you can also click **Save** as soon as the administrative console offers the opportunity to do so.

To discard your changes, click Cancel.

Reviewing the token generation configuration

Review the token generation configuration to make changes or save as needed.

About this task

On the **Summary** tab, review, amend, save, or discard your configuration.

Steps

- To amend your configuration, click the corresponding tab title, then follow the configuration wizard to complete the task.
- To keep your changes, click **Done** and continue with the rest of the configuration.



When editing an existing configuration, you can also click **Save** as soon as the administrative console offers the opportunity to do so.

To discard your changes, click Cancel.

Reviewing the SP STS configuration

Review the service provider (SP) security token service (STS) configuration to make changes or save as needed.

About this task

On the **Summary** tab, review, amend, save, or discard your changes.

Steps

- To amend your configuration, click the corresponding tab title, then follow the configuration wizard to complete the task.
- To keep your changes, click **Done** and continue with the rest of the configuration.



When editing an existing configuration, you can also click **Save** as soon as the administrative console offers the opportunity to do so.

• To discard your changes, click Cancel.

PingFederate Performance Tuning Guide

This section shows you to how to fine-tune a few simple application and system level settings to enable PingFederate to achieve maximum performance of the hardware chosen for your deployment.

The default configuration for PingFederate 10.1 is acceptable for most small size deployments. Missioncritical and high-transaction volume deployments might require additional tuning.

This guide addresses several areas of tuning such as logging, concurrency, memory, and Java-specific tuning options. It is not designed as a one-size-fits-all set of instructions to optimize PingFederate, but more as a checklist of suggestions for areas of the product that can be tuned to improve performance, and any tradeoffs associated with those changes. For ultimate reassurance that any fine-tuned settings will meet your expectations, performance testing in a lab environment is recommended.



Note:

The PingFederate Capacity Planning Guide, is also available as a performance data reference. For more information, see the knowledge base section of support.pingidentity.com.

Logging

This section explains the logging practices of PingFederate and discusses minimizing the system's overall performance impact.

Logging tracks various aspects of the system's overall performance and requires a certain amount of system resources, which affects the system's overall performance. In particular, writing to the log files takes the greatest amount of resources. To minimize the performance impact, PingFederate uses the highperformance asynchronous logger from Log4j 2 for logging runtime and administrative events, including status and error messages used for troubleshooting. To preserve transactional integrity, audit information logs synchronously.

Although the bulk of logging is executed asynchronously, decreasing the amount of information written to log files always provides the best possible performance.

PingFederate only records messages tagged with log level INFO, WARN, ERROR, and FATAL to the server log and the provisioner log. Messages with DEBUG, or TRACE tags, are not recorded to optimize performance. Console logging is also disabled for the same reason.

For troubleshooting purposes, you can adjust the log level to DEBUG in the log4j2.xml file and re-enable console logging.



Important:

When you no longer require debug messages and console logging, turn them off.

Operating system tuning

This section contains tuning recommendations for your operating system.

The tuning recommendations provided here work best in preventing deployment issues in high capacity environments.

Linux tuning

Follow these recommendations for your Linux environment to prevent deployment issues, to increase the performance and capacity of the networking stack, particularly TCP and the file descriptor usage, and to enable PingFederate to handle a high volume of concurrent requests.

Network/TCP tuning

For SystemV, add or modify the following entries in the /etc/sysctl.conf file.

For systemd, you can create a sysctl preload/configuration file in /etc/sysctl.d (for example, 99sysctl.conf) in which to add and modify the following entries.

```
##TCP Tuning##
# Controls the use of TCP syncookies (default is 1)
# and increase the number of outstanding syn requests allowed.
net.ipv4.tcp syncookies=1
net.ipv4.tcp max syn backlog=8192
# Increase number of incoming connections.
# somaxconn defines the number of request sock structures allocated
# per each listen call.
# The queue is persistent through the life of the listen socket.
net.core.somaxconn=4096
# Increase number of incoming connections backlog queue.
# Sets the maximum number of packets, queued on the INPUT side,
# when the interface receives packets faster
# than kernel can process them.
net.core.netdev max backlog=65536
# increase system IP port limits
net.ipv4.ip local port range=2048 65535
# Turn on window scaling which can enlarge the transfer window:
net.ipv4.tcp window scaling=1
# decrease TCP timeout
net.ipv4.tcp fin timeout=10
# Allow reuse of sockets in TIME WAIT state for new connections
# (While this may increase performance, use with caution according
```

```
# to the kernel documentation. This setting should only be enabled
# after the system administrator reviews security considerations.)
net.ipv4.tcp_tw_reuse=1
# Increase the read and write buffer space allocatable
# (minimum size, initial size, and maximum size in bytes)
net.ipv4.tcp rmem = 4096 65536 16777216
net.ipv4.tcp wmem = 4096 65536 16777216
# The maximum number of packets which may be queued
# for each unresolved address by other network layers
net.ipv4.neigh.default.unres qlen=100
net.ipv4.neigh.eth0.unres qlen=100
net.ipv4.neigh.em1.unres qlen=100
# Default Socket Receive and Write Buffer
net.core.rmem default=8388608
net.core.wmem_default=8388608
###############
```

Increase file descriptor limits

Add or modify the following lines in the /etc/security/limits.conf file where *pf_user* is the user account used to run the PingFederate java process or * for all user accounts.

```
pf_user soft nofile 10400
pf_user hard nofile 10400
```

Windows tuning

Follow these recommendations for your Windows environment to prevent deployment issues, to increase the performance and the capacity of the networking stack, specifically the TCP socket, and to enable PingFederate to handle a high volume of concurrent requests.

About this task

Use Command Prompt and a Registry Editor to edit the <code>cmd.exe</code> file and the <code>regedit.exe</code> file respectively.

Steps

- 1. Edit the jvm-memory.options file with a time stamp.
 - a. Start the Command Prompt application cmd.exe.
 - b. Enter netsh int ipv4 show dynamicportrange tcp to view the ephemeral ports.
 - c. Enter netsh int ipv4 set dynamicport tcp start=1025 num=64510 to increase the range of the ephemeral ports using administrative privileges.
 - d. Reboot the server.
 - e. Enter netsh int ipv4 show dynamicportrange tcp to confirm the updated port range.
- 2. Reduce socket TIME_WAIT delay.
 - a. Start the Registry Editor application regedit.exe.
 - b. Go to HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip \Parameters.
 - c. Create a new DWORD 32 bit value and set the name as TcpTimedWaitDelay.
 - d. Set a decimal value of 30.
 - e. Reboot the server.

Concurrency

This section describes how to configure PingFederate to support more concurrent requests to optimize your deployment.

The more requests processed in parallel, the more requests processed over all. Given the appropriate amount of hardware, processing N requests concurrently is typically faster than processing N requests serially.

In PingFederate, there are two main pools of threads that control the level of concurrent user requests: Acceptor Threads and Server Threads. Acceptor threads receive the HTTPS requests and pass those requests on to available server threads to be processed.

Caveats

This topic serves as a guideline for optimizing the concurrency of your deployment. On a large system with multiple CPUs, or cores, a thread pool that is too small will under-utilize the available processor resources. A thread pool that is too large can cause the system to become flooded and unusable.

A good target for the CPU is between 60%-80% utilization when under nominal, standard user load. This way CPU resources are not under-utilized while still allowing room for occasional load spikes. The level of concurrency in PingFederate might need to be decreased, or even increased, depending on the system's configuration, the adapters in use, available memory, and other processes competing for resources. All deployments are different. This section serves as a guideline of where to start when tuning the server.

Tuning the acceptor gueue size

For optimal performance, particularly in larger deployments, PingFederate uses a non-blocking I/O model to process requests. Follow these steps to adjust the queue size parameter for your environment needs.

About this task

As needed, administrators can fine-tune the queue size parameter, acceptQueueSize. If the queue reaches its maximum size, additional requests will receive a connection refused error. If this occurs in your environment, you can increase the value of the acceptQueueSize parameter.

Steps

- 1. Stop PingFederate.
- 2. Edit the <pf install>/pingfederate/etc/jetty-runtime.xml file.



Consider making a backup copy of this file.

3. Go to the following section, and change the acceptQueueSize as needed.

```
<Call id="httpsPrimaryConnector" name="addConnector">
 <Arq>
class="com.pingidentity.appserver.jetty.server.connector.ServerConnector">
      <Arg name="server"><Ref refid="RuntimeServer" /></Arg>
      <Arg name="acceptors" type="int"><Property name="ssl.acceptors"</pre>
default="0"/></Arg>
      <Set name="acceptQueueSize">512</Set>
    </New>
 </Arg>
```

```
</Call>
```

- 4. When finished, save your changes and restart PingFederate.
- 5. For a clustered PingFederate environment, repeat these steps on each engine node as needed.

Tuning the server thread pool

When tuning the server thread pool, set the minimum and maximum number of threads to optimize PingFederate for your needs.

About this task

Tune the server thread pool to a minimum and maximum number to size the system based on expected user load.

Testing has shown that PingFederate performs well when the server thread pool is sized between 25 and 50 server threads per available CPU core, assuming sufficient memory. For more information on managing your memory for PingFederate, see Memory on page 994.



Note:

This guidance might not be advisable for scaling on larger systems. For example, if you are running PingFederate on a system with 24 CPU cores, it does not make sense to size the thread pool at a minimum of 600 threads and a maximum of 1200 unless you expect to normally handle at least 800 concurrent requests.

Steps

- 1. Stop PingFederate.
- 2. Edit the <pf install>/pingfederate/etc/jetty-runtime.xml file.



Tip:

Consider making a backup copy of this file.

3. Go to the following section, and change the thread number as needed.

```
<Get name="ThreadPool">
     <Set name="minThreads" type="int">100</Set><Set name="maxThreads" type="int">200</Set>
     <Set name="detailedDump">false</Set>
```

- 4. Modify the thread number using the following guidelines:
 - Set minThreads to (available CPU cores * 25)
 - Set maxThreads to (available CPU cores * 50)

Example:

If your PingFederate system has 1 CPU with 4 cores, the total available cores is 4. The configuration would be as follows.

```
<Get name="ThreadPool">
    <Set name="minThreads" type="int">100</Set>
    <Set name="maxThreads" type="int">200</Set>
    <Set name="detailedDump">false</Set>
```



Set the minimum number of threads to between 75% and 100% of the number of requests you expect the system to handle most often. Set the maximum to between 25% and 50% higher than the minimum to handle load spikes.

- 5. When finished, save your changes. Restart PingFederate.
- 6. For a clustered PingFederate environment, repeat these steps on each engine node as needed.

Configuring connection pools to datastores

Java Database Connectivity (JDBC) and LDAP datastores use connection pooling to improve the performance and efficiency of communicating with external systems. For optimal performance, a number of connections are required to handle most or all the requests in parallel.

About this task

In the Data & Credential Stores window, set the minimum and maximum values for connection pools to JDBC and LDAP data stores.

Connection pools improve efficiency by maintaining persistent connections to the JDBC or LDAP server preventing the expense of creating the connection on demand. Connection pools also allow more control over the load placed on the back-end server. It might not be necessary to have a connection available for every concurrent request received by the server, but having too few available will cause requests to wait when accessing JDBC and LDAP resources.



Important:

Size the connection pool based on the capacity and limitation of the database or LDAP server. Sizing the connection pool beyond the capability of the back-end server could lead to PingFederate flooding the datastore without any performance improvement. For optimal performance, size connection pools large enough to handle between 50% and 100% of the number of concurrent requests the server is expected to encounter often. For more information on optimizing the connection pool, see Best practices for tuning the JDBC connection pool in the Ping Identity Knowledge Base.

Steps

Choose from configuring connection pools to a JDBC or LDAP datastores:

Datastore type	Configuration steps
Configuring connection pools to JDBC datastores	 a. Go to System # Data & Credential Stores # Data Stores, and select the applicable JDBC datastore. b. Go to Database Config # Advanced. c. On the Advanced Database Options window:
	 Set the Minimum Pool Size value to 50% of the maxThreads value. Set the Maximum Pool Size value to between 75% and 100% of the maxThreads value, subject to the

Datastore type	Configuration steps
	capability of the back-end database server.
	Remember: The maxThreads value is defined in the <pf_install>/pingfederate/etc/jetty-runtime.xml file. For</pf_install>
	more information, see <i>Tuning the server thread pool</i> on page 992).
Configuring connection pools to LDAP datastores	a. Go to System # Data & Credential Stores # Data Stores, and select the applicable LDAP datastore.
	b. Go to the LDAP Configuration # Advanced # Advanced LDAP Options.
	c. Set the Minimum Connections value to 50% of the maxThreads value.
	d. Set the Maximum Connections value to between 75% and 100% of the maxThreads value, subject to the capability of the backend database server.

2. For a clustered PingFederate environment, replicate the changes to all engine nodes on the **System** # **Server** # **Cluster Management** window.

Memory

After the CPU, memory is the most important resource for sizing Java virtual machine (JVM) heap, managing garbage collection, and optimizing the overall performance of your PingFederate deployment.

The Concurrency section describes how to configure PingFederate to support more concurrent requests. This section highlights how supporting increasing concurrency requests can affect PingFederate's performance because these requests require an increase in memory. Because PingFederate is a Java application, it is important to consider how tuning affects, or is affected by, garbage collection. This section is not a guide to garbage collection theory or ergonomics.

JVM heap

The most important tuning for the Java Virtual Machine (JVM) is the size of the heap memory, which ensures adequate memory is available to manage garbage collection and optimize overall performance.

If the demands require more memory than what is currently available, the JVM must grow the heap, if it can, or perform garbage collection to provide memory to allocate. Resizing the heap and garbage collecting can be an expensive processes and negatively impact performance. Sizing the heap to ensure an adequate amount of memory is available but still manageable to garbage collection is important in optimizing overall performance.

PingFederate attempts to optimize JVM heap and garbage collector settings based on available system resources at the time of installation and upgrade. Depending on your environment, you can override these settings at a later time.

Additional considerations

The JVM can grow the heap from the minimum heap variable value up to the maximum heap variable value. However, growing the heap is often an expensive exercise and requests memory from the operating

system. In addition, the JVM must also reorganize the heap to account for the memory being added. To conserve memory in your deployment, set a lower value for the minimum heap than that of the maximum heap to ensure you are not reserving unused memory. If you have enough memory that a certain amount is easily earmarked for the PingFederate server, adjust the size of the heap by setting the minimum heap and maximum heap to the same value. This allows the JVM to reserve its entire heap and decrease the amount of resizing that the JVM needs to perform if the amount of memory in use exceeds the value of the minimum heap.

Garbage collectors

Although the java virtual machine (JVM) configures itself for optimal performance for most situations, it is advantageous to configure it for a specific application.

This section explores additional tuning options to improve the PingFederate service in your deployment. In all cases, test these options to ensure that they do benefit your deployment before enabling them in a production environment.

Parallel collector

By default, the server HotSpot JVM selects the parallel collector on machines with multiple CPUs or cores. On machines with a single CPU, the serial collector is used.

On machines with eight or more hardware threads, the parallel collector uses a fraction of them as the number of garbage collector threads; the fraction is roughly 5/8. If the number of hardware threads drops below 8, the number of garbage collector threads matches the number of hardware threads.

The number of garbage collector threads can be overridden by the -XX:ParallelGCThreads=n JVM option, where n is the desired number of garbage collector threads. It is generally recommended to leave ParallelGCThreads as the default. Specifically, setting to a value greater than the number of CPUs will not improve garbage collection performance as the GC threads will all contend for CPU time, causing the operating system to context switch between them. Setting to less than the number of CPUs can cause longer than necessary pause times because not all available CPU resources will be utilized.

The parallel collection is generational, so minor (young generation only) and major (entire heap) collections do occur. Because PingFederate uses a much larger proportion of short to medium lived objects, consider applying a young generation bias to the JVM heap, which will improve performance because the parallel scavenge copying collector used for the young generation is fast. For more information, see *Young generation bias* on page 996.

Concurrent mark sweep collector

The concurrent mark sweep (CMS) collector suits applications that prefer shorter garbage collection pauses, can afford to share CPU resources with the garbage collector, have a large set of long-lived objects in the tenured (or old) generation, and run on machines with multiple CPUs.

The CMS collection is generational, so minor, young generation only, and major, entire heap, collections occur. Because PingFederate uses a much larger proportion of short to medium lived objects, the CMS collector is not generally the best fit. However, if you use a JVM heap greater than 6 GB on larger machines with eight or more CPUs, the CMS collector can provide shorter pause time than the parallel collector in major collections. Applying a young generation bias to the JVM heap will improve performance when using the CMS collector because it employs the same parallel scavenge copying collector for the young generation as the parallel collector.

Enable the CMS collector using the -XX:+UseConcMarkSweepGC JVM option as needed.

Garbage first collector

The garbage first (G1) collector is best for machines with multiple CPUs and a JVM heap of 6 GB or more. The G1 collector is designed to achieve high throughput while meeting its pause times goal for garbage collection.

Important:

When using the G1 Collector, it is best to remove any sizing options specific to the young generation. The G1 collector self-tunes by adjusting the size and nature of the various heap regions to meet the pause time goal. By setting a fixed amount of memory to be used for young generation regions, it limits its self-tuning options.

When applicable, you can enable the G1 collector using the #XX:+UseG1GC JVM option.

For more information about each garbage collector, see Java Platform, Standard Edition HotSpot Virtual Machine Garbage Collection Tuning Guide from Oracle.

Young generation bias

An overview of the generation memory management model in Java and the processing model of PingFederate, how they create, collect, and manage objects, and how to optimize both for your deployment needs.

In the Java generational memory management model, the young generation creates new objects and the garbage collector gathers them when they are no longer used. For example, if the young generation creates a new object, and a reference is maintained, then the garbage collector moves that object to the old generation.



If you intend to use the garbage first (G1) collector, let the Java virtual machine (JVM) handle this aspect of the heap because specific settings can affect the performance of the G1 collector adversely.

For more information on fine-tuning memory and garbage collection settings and instructions, see Finetuning JVM options on page 1003.

The processing model of PingFederate is mostly geared towards short-lived transactions, such as single sign-on and token processing and not long-held, interactive, user sessions. Because most of the objects created are relatively short-lived, it does not make sense to promote short-lived objects to the old generation because they are not needed for long. When the young generation fills up, space must be made for new objects. To allow for more space, objects still in use in the young generation must be moved to the old generation. Moreover, those objects will become garbage and need to be collected from the old generation.

The old generation is typically more expensive to clean up than the young generation. The old generation is cleaned only during a full garbage collection, which happens when the JVM has almost reached the value of the maximum heap variable -xmx and the entire heap must be cleaned. With multiple threads by default on systems with multiple cores, the young generation is garbage collected more often and so the pauses for collections are shorter.

By default, the JVM tends to size the generations biased to the old generation, giving it most of the space of the heap. This results in moving objects more frequently from the young generation into the old generation to make space for new objects, and more frequent full collections as the old generation fills up. By configuring the JVM to provide more memory to the young generation, it reduces the frequency of full, more costly, collections. You can either specify fixed values for the size of the young generation or modify the ratio of young generation to old generation.

To specify a fixed value for the young generation, use the -XX:NewSize= and -XX:MaxNewSize= arguments. These arguments are to the young generation what the minimum heap variable (-xms) and the maximum heap variable (-Xxx) are to the entire heap. The -XX:NewSize= and -XX:MaxNewSize= arguments define the initial or minimum and the maximum sizes of the young generation. The same reasoning that applies to the minimum and maximum heap variables also applies when adjusting these values.

To specify a ratio between the old and new generation size, use the -XX: NewRatio= argument. For example, setting -XX: NewRatio=3 means that the ratio between the young and old generation is 1:3. The size of the young generation is 25% of the total heap size.

In a mostly short-lived object environment, give 50 - 60% of the heap to the young generation. See the following table for examples.

Setting the JVM options according to young generation sizing

Young generation bias condition	JVM options
To fix a heap size of 2 GB with 50% the young generation bias using the NewSize argument	-Xms2048m -Xmx-2048m -XX:NewSize=1024m -XX:MaxNewSize=1024m
To fix a heap size of 2 GB with 50% the young generation bias using the NewRatio argument	-Xms2048m -Xmx-2048m -XX:NewRatio=1

The memoryoptions utility

Where to find the memoryoptions utility in a PingFederate installation and how the utility's expected behavior differs in a Linux or Windows system.

PingFederate installation and upgrade tools use the memoryoptions utility to detect the available resources at the time of the installation or upgrade and record the recommended options for Java heap and the garbage collector in a configuration file. As needed, administrators can re-run the utility or manually edit the configuration file.

The memoryoptions utility, located in the <pf_install>/pingfederate/bin directory, comes in two variants:

- memoryoptions.bat for Windows
- memoryoptions.sh for Linux

The configuration file, jvm-memory.options, is located in the same bin directory.

Installation and upgrade

When the PingFederate installer is executed for Windows or a subsequent rerun of the memoryoptions utility, it creates a backup copy of the current jvm-memory.options, if any, detects available system resources at the time, and records the recommended options in the jvm-memory.options file. Changes made as a result of the execution of the utility or a manual edit are activated after a restart of PingFederate.

Depending upon the selected tool and whether the jvm-memory.options file exists in the source installation, the expected behavior of the memoryoptions utility differs. In general, the jvm-memory.options file from the source installation is preserved without new recommended values.

memoryoptions and installation

When the PingFederate installer for Windows runs the memoryoptions utility tool or when changes are made from a manual edit, expect the following behaviors to the installation medium.

The PingFederate installer for Windows runs the memoryoptions utility in an attempt to optimize the Java virtual machine (JVM) heap and garbage collector options based on available system resources at the time of installation. As needed, administrators can rerun the utility or manually edit these options at a later time.

PingFederate installation mediums and their expected behaviors from the execution of the memoryoptions utility tool

PingFederate. See the following table for information regarding expected behaviors.

Installation medium	Expected behavior
PingFederate installer for Windows	 The installer creates a new PingFederate installation. The installer runs the memoryoptions utility, which is designed to determine the recommended Java heap and garbage collector options based on the available resources and to record them in the jvm-memory.options file. The installer configures PingFederate to run as a service. The recommended options are activated as the PingFederate service starts.

Installation medium	Expected behavior	
PingFederate product distribution ZIP file	The default jvm-memory.options file becomes part of the new installation as program and default configuration files are extracted from the PingFederate product distribution .zipfile.	
	PingFederate as a console application on Windows <i>or</i> as a console application, or a service on Linux	
	 The JVM options set in the default jvm-memory.options file are activated as PingFederate starts. The default JVM options are conservative. For most deployment scenarios using various physical or virtual resources, run the memoryoptions utility, which is designed to determine the recommended Java heap and garbage collector options based on the available resources and record them in the jvm-memory.options file. As a result of the execution of the memoryoptions utility or a manual edit of the jvm-memory.options file, the JVM options are activated as PingFederate restarts. 	
	PingFederate as a service on Windows	
	• When administrators run the PingFederate service-installation program install-service.bat, located in the <pf_install>/ pingfederate/sbin/win-x86-64 directory, to install the PingFederate Windows service manually, the program runs the memoryoptions utility, which is designed to determine the recommended Java heap and garbage collector options based on the available resources and record them in the jvm- memory.options file.</pf_install>	
	The service-installation program then runs a helper utility generate-wrapper-jvm-options.bat, located in the <pf_install>/pingfederate/sbin/wrapper directory, to read the JVM options from the jvm-memory.options file and create a resource file that the PingFederate Windows service requires to configure its JVM options. The recommended options are activated as the PingFederate service starts.</pf_install>	

memoryoptions and upgrade

Upgrade paths behave differently when changes are executed based on the recommendations and execution of the memoryoptions utility tool.

Depending upon the selected tool and whether the jvm-memory.options file exists in the source installation, the expected behavior of the memory.options utility differs. In general, the jvm-memory.options file from the source installation is preserved without new recommended values. See the following table for information regarding expected behaviors.

Upgrade path

Expected behavior when the jvm-memory.options file does not exist in the source installation

PingFederate installer for Windows

- The installer creates a new PingFederate installation.
- The installer runs the memoryoptions utility, which is designed to determine the recommended Java heap and garbage collector options based on the available resources and to record them in the jvmmemory.options file.
- The installer configures PingFederate to run as a service.
- The recommended options are activated as the PingFederate service starts.

PingFederate Upgrade Utility (upgrade.bat)

The upgrade utility creates a new PingFederate installation based on the source installation and the PingFederate product distribution .zip file. The default jvm-memory.options file becomes part of the new installation as the upgrade utility extracts files from the PingFederate product distribution .zipfile.

PingFederate as a console application on Windows

- The JVM options set in the default jvm-memory.options file are activated as PingFederate starts.
- The default JVM options are conservative. For most deployment scenarios using various physical or virtual resources, run the memoryoptions utility, which is designed to determine the recommended Java heap and garbage collector options based on the available resources and record them in the jvm-memory.options file.
- As a result of the execution of the memoryoptions utility or a manual edit of the jvm-memory.options file, the JVM options are activated as PingFederate restarts.

PingFederate as a service on Windows

• When administrators run the PingFederate service-installation program install-service.bat, located in the <pf_install>/ pingfederate/sbin/win-x86-64 directory, to install the PingFederate Windows service manually, the program runs the memoryoptions utility, which is designed to determine the recommended Java heap and garbage collector options based on the available resources and record them in the jvmmemory.options file.

The service-installation program then runs a helper utility generate-wrapper-jvm-options.bat, located in the $<\!pf_install>\!/\!pingfederate/sbin/wrapper directory$, to read the JVM options from the jvm-memory.options file and create a resource file that the PingFederate Windows service requires to configure its JVM options.

 The recommended options are activated as the PingFederate service starts.

Upgrade path Expected behavior when the jvm-memory.options file exists in the source installation PingFederate installer for The installer creates a new PingFederate installation based on the source Windows installation and copies the jvm-memory.options file from the source installation to the new installation. At the end of the installation, the installer runs the PingFederate serviceinstallation program, which runs a helper utility generate-wrapperjvm-options.bat, located in the <pf install>/pingfederate/ sbin/wrapper directory, to read the JVM options from the jvmmemory.options file and create a resource file that the PingFederate Windows service requires to configure its JVM options. The preserved Java virtual machine (JVM) options are activated as the PingFederate service starts. PingFederate Upgrade The upgrade utility creates a new PingFederate installation based on the Utility (upgrade.sh) source installation and copies the jvm-memory.options file from the source installation to the new installation. The preserved JVM options are activated as the PingFederate service starts.

Restoring the preserved JVM options

When administrators run the PingFederate service-installation program install-service.bat, the upgrade utility creates a new PingFederate installation based on the source installation and copies the jvm-memory.options file from the source installation to the new installation. You can edit the jvmmemory.options file with a time stamp to restore the preserved JVM options that exist in the jvmmemory.options file from the source installation.

About this task

Use the command prompt to edit the jvm-memory.options file with a time stamp.

Steps

- 1. Rename the currentjvm-memory.options file. For example, jvm-memory.options.backup.
- 2. Look for the preserved jvm-memory.options file.



Note:

The preserved file was renamed with a time stamp.



The jvm-memory.options is the file preserved from the source installation.

- 4. Open a command prompt and go to the <pf install>/pingfederate/sbin/wrapper directory.
- 5. Run generate-wrapper-jvm-options.bat.



Note:

This helper utility reads the JVM options from the jvm-memory options file and creates a resource file that the PingFederate Windows service requires to configure its JVM options.

6. Close the command prompt.

The preserved file was renamed with a time stamp.

7. Restart the PingFederate Windows service.

Result

The preserved JVM options are activated as the PingFederate service starts.

Fine-tuning JVM options

Edit the JVM options jvm-memory.options file to customize minimum and maximum heap sizing, garbage collection, and generation specific sizing for your memory use and to optimize PingFederate's performance.

About this task

PingFederate reads java virtual machine (JVM) options from the jvm-memory.options file, located in the <pf install>/pingfederate/bin directory. Any manual modifications or additions should be made in this file.



Remember:

Before making any edits to the file, consider the following:

- Make a backup copy prior to any manual edits.
- The empty lines and comments, indicated by a leading # character, are ignored.
- JVM options do not need a specific organization or order.

Steps

- 1. Edit the <pf install>/pingfederate/bin/jvm-memory.options file.
- 2. To configure a specific heap size, edit the minimum -Xms and maximum -Xmx heap options.

The valid unit qualifiers are k for kilobytes, m for megabytes, and g for gigabytes. In other words, -Xmx1536m and -Xmx1.5g are equivalent.

Example:

For example, to fix the JVM heap size to 2 GB, configure the minimum and maximum options as follows.

```
-Xms2g
```

-Xmx2g

For example, to configure the parallel collector to use four threads, add the following option to the configuration file.

```
-XX:ParallelGCThreads=4
```

- 4. To enable the concurrent mark sweep (CMS) collector, remove the -XX: +UseParallelGC option, or the options pertaining to another garbage collector, and replace it with the -XX: +UseConcMarkSweepGC option.
- 5. To enable the garbage first (G1) collector, remove the -XX:+UseParallelGC option, or the options pertaining to another garbage collector, and replace it with the #XX: +UseG1GC option.



Note:

If you enable the G1 collector, remove any sizing options specific to the young generation. Skip to step 7.

6. To configure young generation-specific sizing options, edit the minimum -XX: NewSize and maximum -XX: MaxNewSize options for the young generation space.

Example:

For example, to fix the young generation bias to 1 GB, set the minimum and maximum options as follows.

```
-XX:NewSize=1024m
-XX:MaxNewSize=1024m
```

- 7. To remove young generation-specific sizing options completely, remove the aforementioned options or add a leading # character.
- 8. To add additional JVM options, insert the applicable options to the file.

Example:

For example, to enable the aggressive options flag, configure the file as follows.

```
. . .
# Enable the aggressive options flag
-XX:+AggressiveOpts
```

The comment is optional.

- 9. When finished, save your changes.
- 10. If PingFederate is configured to run as a service on a Windows server, follow these steps:
 - a. Open command prompt and go to the <pf install>/pingfederate/sbin/wrapper directory.
 - b. Run generate-wrapper-jvm-options.bat.

Result:

This helper utility reads the JVM options from the jvm-memory.options file and creates a resource file that the PingFederate Windows service requires to configure its JVM options.

- c. Close the command prompt.
- 11. Restart PingFederate.
- 12. For a clustered PingFederate environment, repeat these steps on each engine node as needed.

Hardware security modules

When configuring PingFederate to use a hardware security module, be aware of the following performance impact considerations.

You can configure PingFederate to use a hardware security module (HSM) for cryptographic material storage and operations. Standards such as the Federal Information Processing Standard (FIPS) 140-2 require the storage and processing of all keys and certificates on a certified cryptographic module.

For more information on supported configurations for secure material storing and processing, see Supported hardware security modules on page 169.

Performance considerations

Configuring PingFederate to use an HSM for cryptographic material storage and operations can introduce an impact on performance. The level of impact depends on the performance of cryptographic functionality provided by the HSM and the network latency between PingFederate and the HSM. Consult your HSM vendor for performance tuning and optimization recommendations if you plan to use an HSM as part of your PingFederate deployment.

Configuration at scale

You can configure PingFederate to improve the administrative-console experience for your scaling needs.

For deployments that have hundreds of connections or OAuth clients, or both, and observe noticeable delays in the administrative console, administrators can optionally configure PingFederate to create configuration archives during off-peak hours and disable automatic connection validation to improve the administrative-console experience.

References

For more information on memory management and Hotspot Java virtual machine (JVM) arguments for garbage collection tuning, see the following resources.

Memory management

For more information, see Java Platform, Standard Edition HotSpot Virtual Machine Garbage Collection Tuning Guide

from Oracle.

Hotspot JVM arguments

Fore more information, see Java HotSpot VM Options from Oracle.

PingFederate Monitoring Guide

PingFederate provides a range of monitoring options, from simple heartbeat options for checking responsiveness to transaction response-time logging and resource-utilization metrics. These metrics can help you gain insight into the health and performance of your PingFederate deployment.

To help you monitor the performance of a PingFederate deployment, this guide provides the following:

- Suggestions for key performance metrics to monitor and a means by which to monitor them.
- Recommendations about resource-utilization thresholds and patterns.

Monitoring options, including logs that can be used to create Splunk dashboards.

Liveliness and responsiveness

One of the simpler methods for monitoring the performance of a PingFederate deployment involves determining whether the PingFederate Server is available and responsive. To help you identify the status of a server, PingFederate provides a heartbeat request endpoint.

Heartbeat endpoint

If the PingFederate server is running, the process of sending a request to the endpoint /pf/ heartbeat.ping returns an HTTP 200 status. If the request times out or requires an extended amount of time to return, the server might be overloaded or experiencing other difficulties.

If a request requires more than two or three seconds to return, multiple factors in your PingFederate deployment might be responsible. We recommend that you develop a baseline for the desired response time by testing the heartbeat endpoint of your deployment at various times. This endpoint can be useful when load balancing a cluster of PingFederate instances. Some load balancers can alter the number of requests that are sent to a particular server based on the response code received, or the responsiveness of requests that are made to the heartbeat endpoint.

The output of the heartbeat endpoint can be modified to provide performance-related information, such as CPU and memory usage, and response times. The response metrics can help you make better autoscaling decisions. The map size metrics can help you recognize performance issues.

The following example shows a report containing all the PingFederate server metrics available from the heartbeat endpoint.

```
"cpu.load": "6.13",
"total.jvm.memory": "769.13 MB",
"free.jvm.memory": "517.702 MB",
"used.jvm.memory": "251.429 MB",
"total.physical.system.memory": "17.18 GB",
"total.free.physical.system.memory": "358.928 MB",
"total.used.physical.system.memory": "16.821 GB",
"number.of.cpus": "8",
"atm.ref.token.map.size": "99",
"idp.session.registry.session.map.size": "157",
"response.concurrency.statistics.90.percentile": "1",
"response.concurrency.statistics.max": "2",
"response.concurrency.statistics.mean": "1.0093023255813953",
"response.concurrency.statistics.min": "1",
"response.statistics.count": "215",
"response.statistics.window.seconds": "300",
"response.time.statistics.90.percentile": "2",
"response.time.statistics.max": "376",
"response.time.statistics.mean": "144.646511627906977",
"response.time.statistics.min": "101",
"session.state.attribute.map.size": "157",
"sp.session.registry.session.map.size": "0",
"total.failed.transactions": "0",
"total.transactions": "150"
```

The following table describes all the PingFederate server metrics available from the heartbeat endpoint.

Server metrics	Description
cpu.load	Load on the PingFederate server's cores as a percentage of total capacity
total.jvm.memory	Total memory of the JVM

Server metrics	Description
free.jvm.memory	Free memory of the JVM
used.jvm.memory	Used memory of the JVM
total.physical.syst	= Formula
total.free.physical	- Page together that the control of
total.used.physical	L-USYSTSTERMENORY
number.of.cpus	Number of cores on the PingFederate server
atm. <atm>.token.map</atm>	P পর্টার্শিটিer of tokens in the access token manager with the ID specified by <atm></atm>
idp.session.registr	รั บฟริศาธิย์ - อิก id ซิลิท ิเช ริว่าอิ∀ ider sessions
response.concurrenc	ਦੁਸਾਵ ਿ 90th ਸੁਟਾਵਿ r esponse colle urrency (for example, if this value is 124, then 90% of the report samples had response concurrency values below 124)
response.concurrenc	प्रM axiatura hingberest HTTP requests that the PingFederate server processed concurrently
response.concurrenc	ያያለ <mark>ខង់ក ក់រកស់</mark> មាស ሞዋሞ requests that the PingFederate server processed concurrently
response.concurrenc	ะ រូហគ ាកែដក់ទក់ជំណីនិ e r oំr HTTP requests that the PingFederate server processed concurrently
response.statistics	• NONDEr of items considered in the heartbeat report for the time and concurrency statistics
response.statistics	s সান্দর্শনাধন্য প্রাপ্রাপ্রতি conds) for the statistics report (this is an echo of the StatisticsWindowSecs value and provides context for the concurrency and time statistics)
response.time.stati	is Thic ទ០ដើម្បីប្រជាពិក្រជាខិក្សា ទៅ this value is 168, then 90% of the report samples had response times below 168 milliseconds)
response.time.stati	વારા કુલાયા કુલાયા માટે કુલાયા કુલાય
response.time.stati	្នុំ។ ម៉ាច់នក កោលកា milliseconds that the PingFederate server took to respond
response.time.stati	े डींअस्डिसेंग्र ीe in milliseconds that the PingFederate server took to respond
session.state.attri	i ២ឃុប់កាសខាក់ថា ក៏ខ្ញុំកើទិ in the session state
sp.session.registry	୵ ।ବର୍ଗ୍ୟନ୍ତି P ଟିମ୍ୟ ଅଞ୍ୟ ାଠ <mark>ର ଜୁମି ପ</mark> vider sessions
total.failed.transa	^{ឧទាស់} អាមិ្ទ of failed transactions since the PingFederate server started
total.transactions	Number of transactions since the PingFederate server started

The statistics are for a 5 minute interval and they are updated every 30 seconds. The report takes only the first 5000 items into account.

For more information, see Customizing the heartbeat message on page 835

Response-time logging

By default, the audit logs record the processing time for each transaction. With audit logging enabled, you can identify the speed with which PingFederate processes the following transaction types:

- Single sign-on (SSO)
- OAuth
- Security token services (STS)

Depending on your logging configuration, audit logging might not log any transactions. For more information, see *Security audit logging* on page 767.

The following provides examples of the default audit log.

```
2019-11-10 13:24:57,493| tid:cYunBsgybiw_fiRnJjkAhbIXvzc|
AUTHN_SESSION_USED| | 127.0.0.1 | | ac_client| | localhost| IdP| success|
PdFormAdpt| | 17

2019-11-10 13:24:58,720| tid:cYunBsgybiw_fiRnJjkAhbIXvzc| OAuth|
5c60f022-le9d-3fbe-9749-4b9ca5591356| 127.0.0.1 | | ac_client| OAuth20|
localhost| AS| success| PdFormAdpt| | 7
```

Processing times are shown at the end of the entry in milliseconds.

Resource metrics

PingFederate provides mechanisms for obtaining resource metrics including JMX and heartbeat endpoint..

PingFederate provides the following mechanisms for obtaining resource metrics:

- JMX Ping recommends using JMX MBeans because this method provides a more comprehensive set of resource metric counters for analyzing performance. Several tools are available for collecting and analyzing data from JMX MBeans, including many security information and event management (SIEM) tools, like Splunk.
- Heartbeat endpoint For more information about enabling heartbeat message reporting, see Runtime reporting on page 864.

Monitoring discusses the JConsole monitoring tool that is included with the Java SE platform. For more information about the Comprehensive JConsole, see *Troubleshoot with the JConsole Tool* in the Oracle JDK documentation and *The Java Monitoring and Management Console (jconsole)* in the OpenJDK documentation.

Connecting with JMX

You can connect to JMX using local and remote processes.

JConsole permits connections to local and remote Java processes. If your instance of PingFederate is running as a Windows Service, you must connect through the remote option. For more information on connecting to a local process, see *Connecting to a local process* on page 1009. For information on connecting to a remote process, see *Connecting to a remote process* on page 1009.

Connecting to a local process

Unless you are running PingFederate Server as a Windows service, the easiest method by which to launch JConsole on the same machine as the server is to select Local Process.

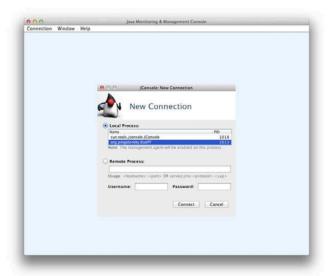
About this task

For information about connecting to a remote process, see Connecting to a remote process on page 1009.

To connect to a local instance and start the monitoring process:

Steps

From the Local Process list, select org.pingidentity.RunPF and then click Connect.





Note:

If you are running the process locally, the system might prompt you to accept the connection as insecure.

Connecting to a remote process

If PingFederate is running as a Windows Service, or if the .org.pingidentity.RunPF class is unavailable in the Local Process list, use this procedure to establish a connection.

About this task

To enable remote JMX monitoring in PingFederate:

Steps

- 1. In the Administrative Console, go to the Security # System Integration # Service Authentication window.
- 2. Define the credentials that are required to connect to the PingFederate JMX service.
- 3. Restart PingFederate to enable the JMX Service.
- 4. In a clustered PingFederate environment:
 - a. Replicate the configuration changes on each node in the cluster.
 - b. Restart each engine node.

5. After you enable the JMX service, connect to the remote JMX service by specifying one of the following:

Choose from:

- The name of the PingFederate server instance.
- The IP address, port 1099 (the default JMX port for PingFederate), and the authentication credentials that the Service Authentication page defines.

Because JMX uses SSL by default when communicating with a remote host, the client host must trust the PingFederate SSL certificate that is presented during setup for JMX. For more information, see Runtime monitoring using JMX on page 865. To disable the use of SSL for JMX, open the /server/default/conf/jmx-remote-config.xml file and set the <item name="jmx.rmi.ssl"> property to false.



Note:

If the JMX client does not trust the JMX certificate, a connection failed SSL message appears.

- 6. If SSL is enabled in jmx-remote-config.xml, import the PingFederate SSL certificate to the client's trusted certificates.
- 7. If SSL if disabled, click **Insecure** to connect.

Monitoring

This topic outlines the key JVM performance metrics for evaluating the performance of a PingFederate deployment.

After a connection is established, you can access the JConsole monitoring interface.

Monitoring clustered PingFederate engines

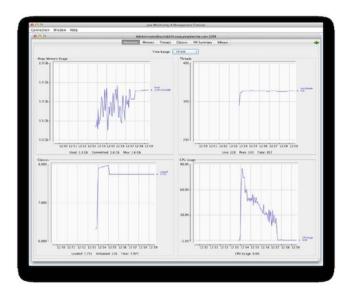
JConsole can be connected to multiple processes. To monitor several instances of PingFederate after a connection is established, click Connection # New Connection and add the additional connection.

Monitoring CPU utilization

The **Overview** tab provides a dashboard of the following performance and resource-utilization charts:

- Heap Memory Usage (cumulative memory that is used by all memory pools).
- Live Threads
- CPU Usage
- Classes (number of classes that are loaded)

This tab provides a high-level view of the JVM's performance metrics.



Use the **Overview** tab to visualize and collect CPU usage data. When your PingFederate deployment is subjected to its normal or expected load, the CPU utilization typically falls between 60 and 80%. If the system registers consistently at 80% or higher, additional CPU resources might be necessary to handle load spikes that occur during peak usage times.

Monitoring memory utilization

The **Overview** tab shows only overall heap usage. To view additional details about memory utilization, click the **Memory** tab, which lets you analyze usage patterns in specific memory pools within the heap. This tab also provides information about the overall heap utilization profile.

Old Generation space

Objects that survive a sufficient number of garbage-collection cycles are promoted to the Old Generation. To view the memory usage in the pool of such objects, click **Memory Pool # PS Old Gen** or **Memory Pool # G1 Old**, depending on the relevant garbage collection. PingFederate services mostly short-lived transactions, like SSO, STS, and OAuth requests and most of the created memory objects are required only for a short period of time.

Although PingFederate makes use of some memory objects that are medium to long lived, such as session data for authentication session, adapter sessions, or single logout functionality, most of the objects that are promoted to the Old Generation are likely to become garbage that requires cleaning up. If the younger generation, or *Eden space*, is not sized appropriately, objects are moved to and retained in the Old Generation before they are collected as garbage. If size limitations prevent the Old Generation from accumulating future garbage as well as longer-lived objects, then garbage-collection cycles occur more frequently.

The Old Generation space is the most important space to monitor. It is easy to identify if the heap is sized and proportioned appropriately for a specific load, based on its usage pattern. The following examples involve two Old Generation usage charts. In both examples, the following examples involve two Old Generation usage charts. In both examples, the same user load executes the same workflow. The size of the heap represents the only difference.

The most important space to monitor. It is easy to identify if the heap is sized andBecause the heap is sized adequately in the first example, memory in the Old Generation rises at a reasonably slow rate. Garbage collection frees around 60 to 75% of the space, and room is available to accommodate the future garbage of newly created objects that are moved from the Eden space, as well as the longer-term objects that remain in use. Although the space is 1 GB in size, the average full (PS MarkSweep or G1 Old Generation) collection time is approximately only 240 milliseconds (0.728 seconds for three collections).

When a heap is sized inadequately, the Old Generation runs out of space. In the following example, the amount of memory that becomes free with each garbage collection shrinks, due to the rate at which objects are promoted from the Eden space.



184 PS MarkSweep (full) collections require garbage collections more frequently, totaling 60 seconds, or an average of 326 milliseconds per collection.

Entire heap space

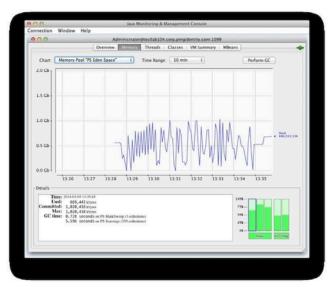
If the heap is sized appropriately for the load that the system must handle, it fills up and is followed by an appreciable drop in usage as a full garbage collection occurs (such as a PS MarkSweep collection triggered by the Old Generation filling up). In this example, the heap rises steadily, with drops from minor collections until a PS MarkSweep collection occurs and collects approximately 70% of the heap.

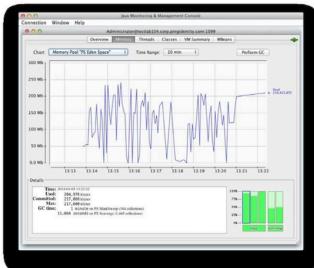
When the heap is undersized, full collections that are performed more frequently return less memory. In the following example, the frequency of JMX data that the JConsole retrieves does not keep pace with the frequency of full collections. As a result, only a fraction of them occur.



Eden space

Regardless of whether the heap is adequately sized or undersized, the usage pattern is nearly identical with the Eden space. This similarity can be due to the sampling frequency of the data-collection tool because the number of samples might be insufficient to show that, with an undersized heap, memory is consumed and subsequently freed with greater frequency. The behavior of garbage collection in the Eden space is such that when it fills, the space is completely emptied by moving live objects to the Survivor and Old Generation spaces. Under load, the pattern resembles a jagged sawtooth, as shown in the following examples of an adequately sized heap and an undersized heap.





Increasing heap size

Because garbage collectors manage memory in the Java Runtime Environment, simply increasing the size of the heap is not always the appropriate solution. The following table outlines the total heap size recommendations for the available garbage collectors, based on available CPU resources. For more information about garbage collectors, see *Garbage collectors* on page 995.

Garbage Collector	Minimum Recommended Number of CPUs	Recommended Heap Size
Parallel	4	6 GB maximum
Concurrent Mark Sweep	12	4 - 6 GB minimum
Garbage First (G1)	12	6 GB minimum

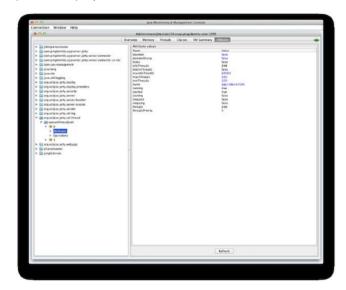
If additional memory is unavailable, or if increasing the size of the heap is inadvisable because of these recommendations, the load that is handled by this instance is probably too high. In such instances, consider adding additional resources to your deployment. To verify whether the load for the instance is too high, check the CPU utilization

To allow for the most efficient management of memory, set the minimum and maximum heap sizes to the maximum allowed values to avoid potentially expensive heap allocation resizing and divide it evenly between the young and old generations. If you are using the Garbage First collector, generational spaces are not specified through command line options because they are managed logically in real time. Even in such instances, we recommend setting the minimum and maximum heap sizes to the maximum allowed values. For more information about fine-tuning the JVM options in the <code>jvm-memory.options</code> file, see Fine-tuning JVM options on page 1003.

Thread pool

The following topic describes the MBeans tab and the recommended number of threads in the pool.

The **MBeans** tab provides access to the JMX Managed Beans and their available attributes and operations. Of particular interest are the <code>queuedthreadpool</code> instances that are available within the <code>org.eclipse.jetty.util.thread</code> bean. For example, instance 0 represents the thread pool that handles runtime requests. When you click the **Attributes** item for instance 0, the current state of the thread pool is displayed.



The number of threads in the pool (the threads attribute) can be compared to the number of threads that are not currently in use (the idleThreads attribute). Ideally, a sufficient number of threads is available to handle load spikes that occur during peak usage times, while also limiting the number of idle threads that are running. If the thread pool is too small, requests might be blocked. If the thread pool is too large, memory might be used unnecessarily, and CPU contention might increase, limiting the processing effectiveness.

We recommend allowing for 10 to 25% more threads than are typically active while the system executes a normal or expected load. Because most of your users will not be active at the same time, set the minimum threads to 10% above the average number of active threads that are observed during monitoring.

We also recommend setting the maximum number of threads to 25% above the average number of active threads that are observed while monitoring under expected load conditions. Make certain to weigh this recommendation against the observed CPU utilization metrics and the suggestions in the *Performance Tuning Guide*.

This section provides an overview of the available logging, reporting, and troubleshooting features for PingFederate.

PingFederate Logs

The server.log file represents the primary troubleshooting log. Along with an HTTP trace from the browser, which can be generated from a debugging application like Fiddler, this file is helpful for identifying issues that must be resolved. The following table identifies the available PingFederate logs and summarizes their purposes.

Name	Purpose
admin.log	Records the actions that users of the Administrative Console perform.
admin-event-detail.log	If detailed event logging is enabled, this log records detailed information about each applicable administrative-console event that users of the Administrative Console perform.
admin-api.log	Records the actions that users of the administrative API perform.
runtime-api.log	Records the actions that API users perform by using the OAuth Client Management Service, the OAuth Access Grant Management Service, and the Session Revocation API.
transaction.log	Records individual identity-federation runtime transactions at specified levels of detail.
audit.log	Records a selected, configurable subset of transaction log information plus additional details. Intended for security-audit and regulatory-compliance purposes.
provisioner-audit.log	Records outbound provisioning events intended for security-audit purposes.
provisioner.log	Records provisioning activity only. Useful when troubleshooting issues that relate to provisioning.
server.log	Records PingFederate runtime and administrative server activities. For more information about the primary troubleshooting log, see <i>Creating an error-only server log</i> .
init.log	Records only Jetty messages that are generated prior to starting PingFederate.

Creating an error-only server log

This section describes am approach for modifying your log4j2.xml file, which can be sent to a security information and event management (SIEM) tool, such as Splunk. You can configure alerts to send notifications when such events occur, or to improve the monitoring of these events.

About this task

We recommend using the server.log file for error-level messages. Even when levels are down to a minimum, the server log generates large amounts of information in an active production environment. As an alternative, you can set up a specific log to log only ERROR and higher.

To change your log4j2.xml file to enable a separate log file:

Steps

1. Create an appender.

The easiest way to create an appender is to copy an existing one as a base. In the following example, the RollingFile is the same one that the server.log file uses. Bold text identifies items that have been changed.

```
<!-- Error Only Main Log: A size based file rolling appender -->
<RollingFile name="FILEERR" fileName="${sys:pf.log.dir}/server.error.log"</pre>
          filePattern="${sys:pf.log.dir}/server.error.log.%i"
ignoreExceptions="false">
<PatternLayout>
     <!-- Uncomment this if you want to use UTF-8 encoding instead
         of system's default encoding.
     <charset>UTF-8</charset> -->
     <pattern>%d %X{trackingid} %-5p [%c] %m%n</pattern>
</PatternLayout>
<Policies>
     <SizeBasedTriggeringPolicy</pre>
             size="10000 KB" />
</Policies>
<DefaultRolloverStrategy max="5" />
</RollingFile>
```

2. At the end of your log4j2.xml file, set the appender that you created in the previous step for AsyncRoot.

In this example, the level attribute indicates the level of messages that are sent to the log file.

3. Remove the attribute additivity="false" from all other loggers that contain a reference to the File appender.

```
Logger name="org.sourceid.saml20.util.SystemUtil" level="INFO"
additivity="false">
  <!--<AppenderRef ref="CONSOLE" /> -->
  <AppenderRef ref="FILE" />
```

Becomes:

```
<Logger name="org.sourceid.saml20.util.SystemUtil" level="INFO" >
<!--<AppenderRef ref="CONSOLE" /> -->
```

<AppenderRef ref="FILE" />

4. Make this change on all nodes within the cluster.



To expedite this step, we recommend creating a base file with the appropriate changes and copying it to all the nodes.

5. Restart PingFederate.

Splunk dashboards and audit logs

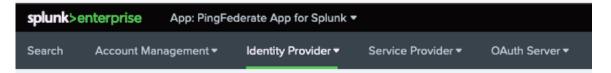
Ping provides a free Splunk for PingFederate application that customers can use to create dashboards.

This application takes advantage of the *Writing audit log for Splunk* on page 777 and *Outbound provisioning audit logging* on page 771, which can be enabled in log4j2.xml file.

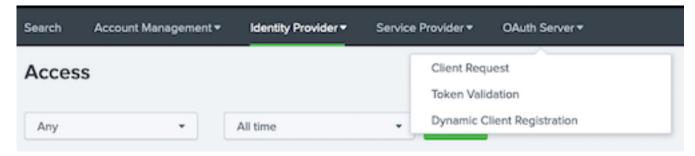
Examples of Splunk dashboards

To help you review different events, the following dashboards are available from the top-level menu of the PingFederate app for Splunk:

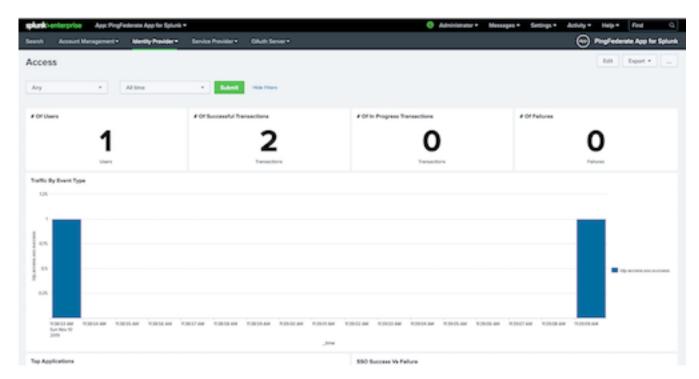
- Account Manager
- Identity Provider
- Service Provider
- OAuth Server



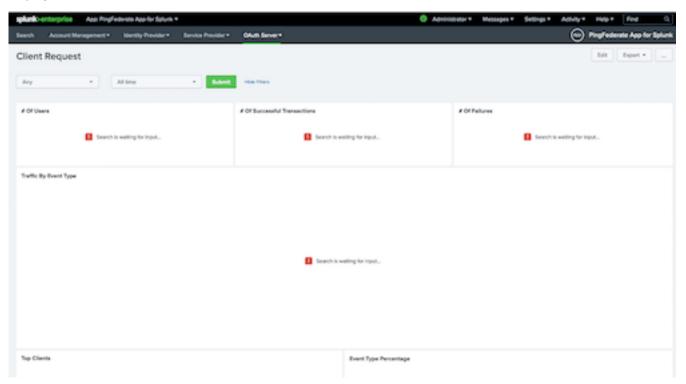
Click a menu item to view its sub-menus, as the following example shows for **OAuth Server**.



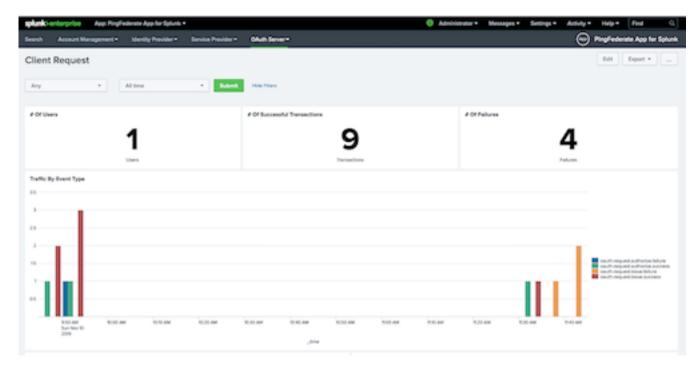
The following image shows the **Identity Provider Access** sub-menu dashboard with examples from the security audit log entries.



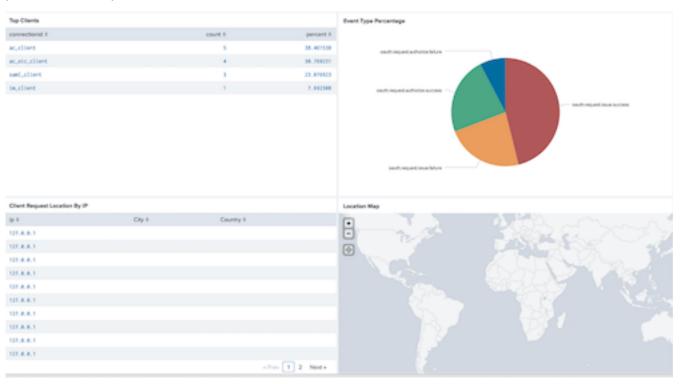
After you select a sub-menu, an image like the following **OAuth Server Client Request** example is displayed while the dashboard waits for the search results.



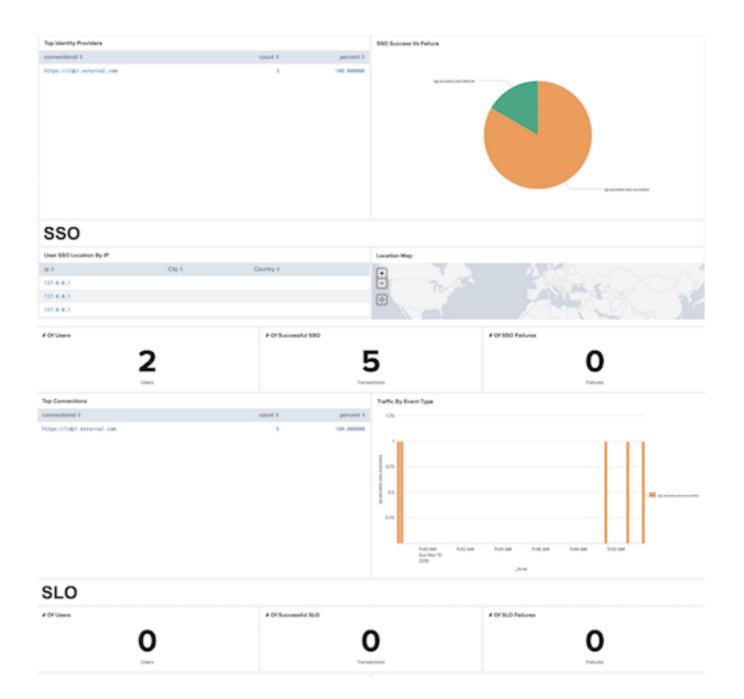
After you click **Submit**, the dashboard displays the following search results for the client request.



To view additional results, scroll downward or select another page. The following **Client Request** page provides an example.



The following images provide additional examples of the **Service Provider Access** sub-menu dashboard.



SDK Developer's Guide

The PingFederate SDK enables integration with identity providers (IdPs) and service providers (SPs). The SDK allows developers to build custom implementations for communicating authentication and security information between PingFederate and their enterprise environment.

Intended audience

This guide is for application developers and system administrators responsible for extending PingFederate, including development of:

- Authentication adapters to integrate web applications or identity-management systems
- Authentication selectors to direct single sign-on (SSO) authentication to instances of authentication adapters based on specified conditions

- WS-Trust Security Token Service (STS) token translators, including token processors and token generators
- Custom data source drivers
- Password credential validators
- Identity store provisioners
- Notification publishers

Introduction

The PingFederate Java SDK consists of several APIs, including:

- Adapter and STS token-translator interfaces
- Authentication selector interfaces
- Custom data source interfaces
- Password credential validator interfaces
- Identity store provisioner interfaces
- Notification publisher interface

These interfaces allow users to create their own custom PingFederate plugins to suit their organization's needs. This SDK provides a means to develop, compile, and deploy custom plugins to PingFederate. The package also contains example plugins for reference. These example plugin projects are located in the <pf install>/sdk/plugin-src directory.

The PingFederate SSO integration overview on page 60 describes the prebuilt authentication adapters Ping Identity provides for integrating web applications and identity-management systems with PingFederate. Review this document before building your own adapter to see if an available adapter fits your use case.



Important:

Custom components might not work the same way after upgrading PingFederate. When upgrading, ensure you thoroughly retest the behavior of customizations in a non-critical upgraded environment.

Adapter and STS token-translator interfaces

The adapter and token-translator APIs enable PingFederate integration with IdPs or SPs. In addition to providing requisite runtime integration, adapter token-translator APIs are configurable UI plugins that allow you to render custom configuration windows.



Note:

Suitable adapter or token-translator implementations for your deployment might already exist. Before developing your own custom solution, see the Ping Identity Downloads website for available implementations.

Authentication selector interfaces

Authentication selectors provide a mechanism to choose among multiple authentication sources and to direct a user to use a particular adapter or IdP connections. For example, an authentication selector might map internal corporate users to use one adapter and map external non-corporate users to a different adapter. Authentication selectors are configurable UI plugins that allow you to render custom configuration windows.

Custom data source interfaces

The custom data source API is a set of Java interfaces that enable PingFederate to integrate with data stores not covered by existing Java Database Connectivity (JDBC) or LDAP drivers. This allows developers to retrieve attributes from their choice of data source during attribute fulfillment. Custom data source interfaces are configurable UI plugins that allow you to render custom configuration windows.

Password credential validator interfaces

The password credential validator interfaces allow developers to define credential validators that verify a given username and password in various contexts throughout the system. For example, credential validators are used to configure OAuth Resource Owner authorization grants and the HTML Form Adapter.

Identity store provisioner interfaces

Identity store provisioners provide a mechanism for provisioning and deprovisioning users to external user stores. For example, you can configure a custom identity store provisioner within an inbound provisioning IdP connection to provision users using the System for Cross-domain Identity Management (SCIM) protocol. Identity store provisioners are configurable UI plugins which allow you to render custom configuration windows.

Notification publisher interface

PingFederate delivers messages to administrators and end users based on notification publisher settings. Developers can implement custom notification publishers using the NotificationPublisherPlugin interface.

Ping Identity Global Client Services

If you need assistance in using the SDK, visit the Ping Identity Support website.

Additional Documentation

- Javadocs provide detailed reference information for developers. The Javadocs are located in the f install/pingfederate/sdk/doc directory.
- The PingFederate *Administrator's Reference Guide* on page 214 provides background information and UI configuration details.
- The user guides for Java, .NET, and PHP integration kits show examples of SDK implementations.

SDK directory structure

This topic describes the directory and build components that comprise the SDK.

The PingFederate SDK directory (<pf install>/pingfederate/sdk) contains the following:

- plugin-src/ The directory where you place your custom plugin projects. This directory also
 contains example plugin implementations showing a wide range of functionality. You can use these
 examples to develop your own implementations.
- doc/ Contains the SDK Javadocs. Open index.html to get started.
- lib/ Contains libraries used for compiling and deploying custom components into PingFederate.
- build.properties Contains properties used by the Ant build script, build.xml, to compile and deploy your custom components. Do not modify this file. If you need to override a property, use build.local.properties.
- build.local.properties Allows you to specify which project you want to build and define properties specific to your environment. Use this file to declare the project you want to build.
- build.xml The Ant build script used to compile, build, and deploy your component. This file should not need modification.

The Java SDK, along with Apache Ant, enables you to create directories for your project and use the build script to build, clean, or deploy it. For more information, see *Developing your own plugin* on page 1024.

Developing your own plugin

Set up a development environment within the SDK and use it to create a plugin.

Before you begin

Ensure you have the Java SDK and Apache Ant installed.

About this task

The Java SDK, along with Apache Ant, enables you to create directories for your project and use the build script to build, clean, or deploy it.

Steps

- 1. Create a new project directory in the <pf install>/pingfederate/sdk/plugin-src directory.
- 2. In the new project directory, create a subdirectory named java. This is where you place the Java source code for your implementation. Follow standard Java package and directory structure layout.
- 3. If your project depends on third-party libraries, create another subdirectory called lib and place the necessary .jar files in it.
- 4. Edit the build.local.properties file and set target-plugin-name to specify the name of the directory that contains your project.
- 5. Run the appropriate target to clean, build, or deploy your plugin.

To display a list of available build targets, run ant from <pf install>/pingfederate/sdk.

```
[java] Main targets:
[java]
[java] clean-plugin Clean the plug-in build directory
[java] deploy-plugin Deploy the plug-in jar and libs to PingFederate
                  Package the plug-in jar
[java] jar-plugin
[java]
[java] Default target: help
```



Note:

Because it packages the .jar files with additional metadata to make them discoverable by PingFederate, we recommend building the project with the build.xml file included in the SDK.

Implementation guidelines

The following topics provide programming guidance for developing custom interfaces.

- Shared plugin interfaces on page 1025
- Developing IdP adapters on page 1026
- Developing SP adapters on page 1029
- Developing token processors on page 1030
- Developing token generators on page 1031
- Developing authentication selectors on page 1031
- Developing data source connectors on page 1032

- Developing password credential validators on page 1034
- Developing identity store provisioners on page 1034
- Developing notification publishers on page 1040
- Building and deploying your project

For more details about interfaces discussed here and additional functionality, see the SDK Javadocs.

Shared plugin interfaces

Plugin implementations generally invoke methods categorized as either configurable or describable. This document describes these types of plugins and how they are used in PingFederate.

Configurable plugin

Any custom plugin that requires UI settings is configurable and implements the <code>ConfigurablePlugin</code> interface. This ensures that PingFederate loads the plugin instance with the correct configuration settings.

All plugin types implement the ConfigurablePlugin interface and must define the following within the ConfigurablePlugin interface to enable configuration loading.

```
void configure(Configuration configuration)
```

During processing of a configurable plugin instance, PingFederate calls the ConfigurablePlugin.configure() method and passes a Configuration object. The Configuration object provides the plugin adapter instance configuration set by an administrator in the PingFederate UI.

The SpAuthnAdapterExample.java sample provided with the SDK shows how to use this method to initialize an adapter instance from a saved configuration. After your implementation loads the configuration values, the plugin instance can use them in other method calls.

Describable plugin

Any plugin that requires configuration windows in the PingFederate administrative console is a describable plugin. Most plugins implement the <code>DescribablePlugin</code> interface to ensure that PingFederate renders the correct UI components based on the returned <code>PluginDescriptor</code>.

Adapter and custom data source plugins are special cases and do not implement the DescribablePlugin interface. However, they still return a plugin descriptor (AuthnAdapterDescriptor and SourceDescriptor) and are still describable plugins.

All describable plugins must define a UI descriptor. Use one of the following methods to implement a UI descriptor, depending on the type of plugin:

• For plugins using the DescribablePlugin interface

```
PluginDescriptor getPluginDescriptor()
```

• For adapter plugins

```
AuthnAdapterDescriptor getAdapterDescriptor()
```

For custom data source plugins

```
SourceDescriptor getSourceDescriptor()
```

Describable plugins can return a subclass of PluginDescriptor, so the return type might differ between plugin implementations. Your plugin implementation populates PluginDescriptor with FieldDescriptors, FieldValidators, and Actions and is presented as a set of UI components in the PingFederate administrative console.



Some plugin types offer concrete descriptor implementations for developers. The Javadocs and examples provided with the SDK show which descriptor classes are available for each plugin type. The examples also show you how to use FieldDescriptors, FieldValidators, and Actions to define your plugin descriptor.

Developing IdP adapters

PingFederate uses a select few of your identity provider's (IdP's) methods to call requests, handle logout requests, and perform other functions.

IdP authentication adapter interface

Create an IdP adapter by implementing the IdpAuthenticationAdapterV2 interface. Implementing this interface requires the following Java packages:

- org.sourceid.saml20.adapter.idp.authn
- org.sourceid.saml20.adapter.gui
- org.sourceid.saml20.adapter.conf

For each IdP adapter implementation, in addition to the methods described under Shared plugin interfaces on page 1025, you must define the following:

- Session lookup
- Session logout

IdP adapter session lookup

PingFederate invokes the lookupAuthN() method of your IdP adapter to look up user-session information to handle a request, regardless of whether the request is for IdP- or service provider (SP)initiated single sign-on (SSO), an OAuth transaction, or direct IdP-to-SP adapter processing.



Note:

The IdentityStoreUserProvisioner interface is deprecated. Developers should implement either the IdentityStoreProvisionerWithFiltering or IdentityStoreProvisioner interfaces.

```
java.util.Map lookupAuthN(javax.servlet.http.HttpServletRequest req,
            javax.servlet.http.HttpServletResponse resp,
            java.lang.String partnerSpEntityId,
            AuthnPolicy authnPolicy,
            java.lang.String resumePath)
            throws AuthnAdapterException, java.io.IOException
```



Note:

The IdpAuthenticationAdapterV2 interface provides an overloaded version of the lookupAuthN() method. The other overloaded lookupAuthN() version is deprecated; you should only use the one within the IdpAuthenticationAdapterV2 interface. Use this interface if your adapter requires additional parameters from PingFederate. For a complete list of available parameters, see the IdpAuthenticationAdapterV2 interface in the Javadocs.

In most implementations, a user's session information or a reference to it is communicated to PingFederate using the HttpServletRequest parameter, which is passed to the lookupAuthN() method. For

example, the implementation can pass the user's session information from the IdP application as a cookie or query parameter.

If the request from the user's browser does not contain the necessary information to identify the user, you can use the HttpServletResponse parameter in various ways to retrieve the user's session data, such as creating a 302 redirect or presenting a web page asking for credentials. If your adapter implementation uses the HttpServletResponse parameter to retrieve the user's session, you must, return the user's browser to the URL declared in the resumePath parameter that was set by the PingFederate runtime server. The resumePath parameter is a relative URL signaling PingFederate that a user is continuing an already initiated SSO transaction.



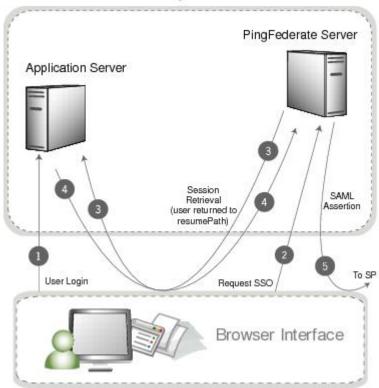
When creating a custom adapter, you can design it to render a template for processing and returning HTML to the user's browser using the TemplateRendererUtil class. A sample (template-renderadapter-example) is included in the sdk/plugin-src directory of your PingFederate instance.

If your adapter implementation writes to the HttpServletResponse to retrieve the user's session data. we recommend returning the URL specified by the resumePath parameter at all times, whether the retrieval succeeds or fails. This ensures the adapter does not interrupt the adapter chain if it is used with the composite adapter. The composite adapter allows an administrator to chain together a selection of available adapter instances for a connection. At runtime, adapter chaining means that SSO requests are passed sequentially through each adapter instance until one or more authentication results are found for the user. If the browser is unable to return to the resumePath URL at all times, then it could interrupt the adapter chain, causing unexpected results for the composite adapter.

For some authentication mechanisms, not all adapters can return the browser to the resumePath URL. Do not use such adapters with the composite adapter's "Sufficient" chaining policy. For more information, see Composite Adapter.

The following diagram illustrates the request sequence of an IdP-initiated SSO scenario that uses the URL specified by the resumePath parameter.

Identity Provider



Processing steps

- 1. User logs in to a local application or domain through an authentication mechanism such as an identity-management system.
- 2. User requests access to a protected resource located in the service provider (SP) domain. The link or other mechanism invokes the PingFederate SSO service.
- 3. PingFederate invokes the designated adapter's lookup method, including the resumePath parameter.
- 4. The application server returns the session information and redirects the browser along with the returned information to the resumePath URL.
- 5. PingFederate generates a SAML assertion and sends the browser with the SAML assertion to the SP's SAML gateway.

IdP adapter session logout

During single logout (SLO) request processing, PingFederate invokes your IdP adapter's logoutAuthN() method to terminate a user's session. This method is invoked during IdP- or SP-initiated SLO requests.

```
boolean logoutAuthN(java.util.Map authnIdentifiers, javax.servlet.http.HttpServletRequest req, javax.servlet.http.HttpServletResponse resp, java.lang.String resumePath) throws AuthnAdapterException, java.io.IOException
```

Like the <code>lookupAuthN()</code> method, the <code>logoutAuthN()</code> method has access to the user's <code>HttpServletRequest</code> and <code>HttpServletResponse</code> objects. Use these objects to retrieve data about the user's session and to redirect the browser to an endpoint used to terminate the session at the application. The <code>resumePath</code> parameter contains the URL to which the user is redirected to complete the SLO process.

Developing SP adapters

This topic describes how to create a service provider (SP) adapter, as well as the methods used during SP session creation, SP adapter session logout, and SP account linking.

SP authentication adapter interface

Create service provider (SP) adapters by implementing the SPAuthenticationAdapter interface. Implementing this interface requires the following Java packages:

```
org.sourceid.saml20.adapter.sp.authnorg.sourceid.saml20.adapter.guiorg.sourceid.saml20.adapter.conf
```

For each SP adapter implementation, in addition to the methods described in *Shared plugin interfaces* on page 1025, you must define:

- SP session creation
- SP adapter session logout
- SP account linking

SP session creation

PingFederate invokes the createAuthN() method during the processing of a single sign-on (SSO) request to establish a security context in the external application for the user.

```
java.io.Serializable createAuthN(SsoContext ssoContext,
  javax.servlet.http.HttpServletRequest req,
  javax.servlet.http.HttpServletResponse resp,
  java.lang.String resumePath)
```

This method resembles the IdpAuthenticationAdapter.lookupAuthN() method in terms of the objects passed to it and its support for asynchronous requests using the HttpServletResponse and resumePath parameters. It also accepts an SsoContext object, which has access to information such as user attributes and the target destination URL.

SP adapter session logout

PingFederate invokes the logoutAuthN() method during a single logout (SLO) request to terminate a user's session with the external application.

```
boolean logoutAuthN (java.io.Serializable authnBean, javax.servlet.http.HttpServletRequest req, javax.servlet.http.HttpServletResponse resp, java.lang.String resumePath) throws AuthnAdapterException, java.io.IOException
```

The HttpServletResponse and resumePath objects are available to support scenarios where the user's browser redirects to an additional service to clean up any remaining sessions.

SP account linking

PingFederate invokes the <code>lookupLocalUserId()</code> method during an SSO request when the identity provider (IdP) connection uses account linking but no account link for this user is yet established.

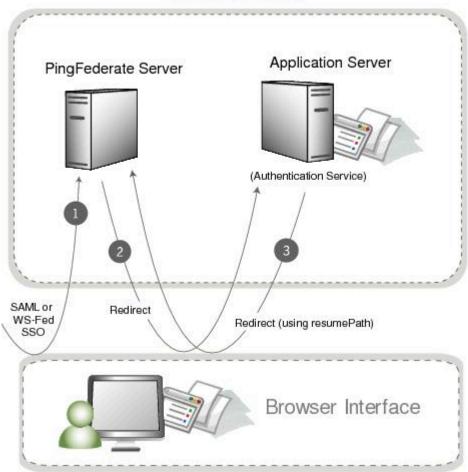
```
java.lang.String lookupLocalUserId(
    javax.servlet.http.HttpServletRequest req,
    javax.servlet.http.HttpServletResponse resp,
    java.lang.String partnerIdpEntityId,
    java.lang.String resumePath)
```

After the account link is set, PingFederate maintains this information until the user defederates, which occurs when the user clicks a hyperlink redirecting them to the /sp/defederate.ping PingFederate endpoint.

The HttpServletResponse and resumePath objects are used to send the user to a local service where the user authenticates. After authentication, the user is redirected to the URL specified in the resumePath parameter and PingFederate completes the account link.

The following diagram illustrates a typical account-link sequence.





Use the ${\tt HttpServletRequest}$ to read a local session token. The <code>lookupLocalUserId()</code> method should return a local user identifier ${\tt String}$ object.

Developing token processors

Create a token processor by implementing the TokenProcessor interface.

The following Java packages are required for implementing the TokenProcessor interface:

- org.sourceid.saml20.adapter.attribute
- org.sourceid.saml20.adapter.idp.authn
- org.sourceid.saml20.adapter.gui
- org.sourceid.saml20.adapter.conf
- org.sourceid.wstrust.model
- org.sourceid.wstrust.plugin

- org.sourceid.wstrust.plugin.process
- com.pingidentity.sdk

For each token-processor implementation, in addition to the methods described under *Shared plugin interfaces* on page 1025, you must define the TokenContext processToken(T token) method.

PingFederate invokes the processToken() method when processing a security token service (STS) request to perform necessary operations for determining the validity of a token. The type parameter T must extend, at a minimum, the type SecurityToken. The type BinarySecurityToken is also available to represent custom security tokens that can be transported as Base64-encoded data.

Developing token generators

Create a token-generator implementation by using the TokenGenerator interface..

The following Java packages are required for implementing the TokenGenerator interface:

- org.sourceid.saml20.adapter.sp.authn
- org.sourceid.saml20.adapter.gui
- org.sourceid.saml20.adapter.conf
- org.sourceid.wstrust.model
- org.sourceid.wstrust.plugin
- org.sourceid.wstrust.plugin.process
- com.pingidentity.sdk

For each token-generator implementation, in addition to the methods described under *Shared plugin interfaces* on page 1025, you must define the SecurityToken generateToken (TokenContext attributeContext) method.

PingFederate invokes the <code>generateToken()</code> method when processing a security token service (STS) request to perform necessary operations for generation of a security token. The type <code>BinarySecurityToken</code> is available and you can use it to represent custom security tokens that can be transported as <code>Base64-encoded</code> data. The <code>TokenContext</code> contains subject data available for insertion into the generated security token.

Developing authentication selectors

This topic describes aspects of authentication selectors within the context of PingFederate, including implementation, context selection, and callbacks.

Authentication selector interface

Authentication selectors allow PingFederate to choose an appropriate authentication source, an identity provider (IdP) adapter or an IdP connection (for federation hub use cases), based on criteria defined in the authentication selector instance.

When creating an authentication selector, use the following primary Java packages:

- org.sourceid.saml20.adapter.gui
- org.sourceid.saml20.adapter.conf
- com.pingidentity.sdk

For each authentication selector implementation, in addition to the methods described under *Shared plugin interfaces* on page 1025, you must define the following:

- Context Selection
- Authentication selector callback

Context selection

PingFederate calls the selectContext() method to determine which authentication source to select. The mappedAuthnSourcesNames contains the list of AuthenticationSourceKeys and names that are available for the selector to reference.

```
AuthenticationSelectorContext selectContext(HttpServletRequest req,
 HttpServletResponse resp,
Map<AuthenticationSourceKey, String> mappedAuthnSourcesNames,
 Map<String, Object> extraParameters,
 String resumePath)
```

The HttpServletRequest can evaluate cookies, parameters, headers, and other request information to determine which authentication source to select. The HttpServletResponse also helps determine the appropriate authentication source to select if the authentication selector requires user interaction. If the resp object is written to, it is considered a committed response and returned to the user's browser. The resumePath is a relative URL used in conjunction with the resp object, such that the user's browser is sent to this URL to resume the single sign-on (SSO) work flow.

Once an authentication source is selected, you can create an AuthenticationSelectorContext to denote which authentication source to use. You can reference the selected authentication source by its ID or by its context, which is a name that decouples authentication selectors from the configured IDs.

Authentication selector callback

PingFederate calls the callback() method after authenticating against a selected source. The callback() method allows authentication selectors to update resulting attributes, set cookies, or perform other custom functions.

```
void callback(HttpServletRequest req,
HttpServletResponse resp,
Map authnIdentifiers,
AuthenticationSourceKey authenticationSourceKey,
AuthenticationSelectorContext authnSelectorContext);
```



Note:

Writing content to the resp object in the callback() method is not supported, and doing so might result in unexpected behavior. Setting cookies is supported.

Developing data source connectors

Use PingFederate to query various data sources or build data source connectors to process customized data sources.

PingFederate can query data sources for a variety of purposes using LDAP or Java Database Connectivity (JDBC) interfaces. Use the PingFederate SDK to build data source connectors to query additional data source types. Examples of other data sources include a web service, a flat file, or a different way of using a JDBC or LDAP connection than what is supplied by PingFederate.

The following are the primary Java packages used to build a custom data source:

- com.pingidentity.sources
- com.pingidentity.sources.gui

For each implementation described in Shared plugin interfaces on page 1025, you must define the following:

- Connection testing
- Available fields retrieval

Data source query handling

Data source connection testing

```
boolean testConnection()
```

When associating a custom data source with an identity provider (IdP) or service provider (SP) connection, PingFederate tests connectivity to the data source by calling the testConnection() method. Your implementation of this method should perform the necessary steps to demonstrate a successful connection and return true, or return false if your implementation cannot communicate with the datastore. A false result prevents an administrator from continuing with the data source configuration.

Data source available fields retrieval

```
java.util.List<java.lang.String> getAvailableFields()
```

PingFederate calls the getAvailableFields () method to determine the available fields that can be returned from a query of this data source. These fields are displayed to the PingFederate administrator during the configuration of a data source lookup, and the administrator selects the attributes from the data source and maps them to the adapter or attribute contract. PingFederate requires at least one field returned from this method.

Data source query handling

```
java.util.Map<java.lang.String,java.lang.Object> retrieveValues(
  java.util.Collection<java.lang.String> attributeNamesToFill,
  SimpleFieldList filterConfiguration)
```

When processing a connection using a custom data source, PingFederate calls the retrieveValues() method to perform the actual query for user attributes. This method receives a list of attribute names populated with data. The method can also receive a filterConfiguration object populated with a list of fields. Each field contains a name/value pair determined at runtime and collectively used as the criteria for selecting a specific record. In most cases, the criteria are used to locate additional user attributes.

Create the filter criteria selections needed for this lookup by passing back a CustomDataSourceDriverDescriptor, an implementation of SourceDescriptor, from the getSourceDescriptor() method. A CustomDataSourceDriverDescriptor can include a FilterFieldDataDescriptor composed of a list of fields that can be used as the query criteria. This list of fields is displayed similarly to the other UI-descriptor display fields.



The filterConfiguration object is set and populated with a list of fields only if the data source was defined with a CustomDataSourceDriverDescriptor. If the CustomDataSourceDriverDescriptor was not used in the definition of the data source, the filterConfiguration object is set to null.



Important:

To pass runtime attribute values to the filter, an administrator must reference the attributes using the \${attribute name} format when defining a filter in the PingFederate administrative console.

Once all relevant attributes are retrieved from the data source, they must be returned as a map of name/ value pairs, where the names correspond to the initial collection of attribute names passed into the method and the values are the attributes.

Developing password credential validators

Password credential validators allow PingFederate administrators to define a centralized location for username/password validation, allowing PingFederate configurations to reference validator instances.

To implement a custom password credential validator, import the following Java packages:

- org.sourceid.saml20.adapter.gui
- org.sourceid.saml20.adapter.conf
- org.sourceid.util.log
- com.pingidentity.sdk
- com.pingidentity.sdk.password

For each implementation, in addition to the methods described under Shared plugin interfaces on page 1025, you must define the following method.

```
AttributeMap processPasswordCredential(String username,
  String password)
  throws PasswordValidationException
```

This method takes a username and password and verifies the credential against an external source. If the credentials are valid, it returns an AttributeMap containing at least one entry representing the principal. If the credentials are invalid, then it returns null or an empty map. If the plugin was unable to validate the credentials (for example, due to an offline host or network problems), it returns a PasswordValidationException.

To enable password changes in a password credential validator, implement the com.pingidentity.sdk.password.ChangeablePasswordCredential interface.

To enable password resets in a password credential validator, implement the com.pingidentity.sdk.password.ResettablePasswordCredential interface.



Note:

Depending on your password management system, you might need additional system configuration to enable password changes. For example, you can change passwords in Active Directory only if LDAPS is enabled.

Developing identity store provisioners

Create an identity store provisioner by implementing either the

IdentityStoreProvisionerWithFiltering Or IdentityStoreProvisioner interface.

Both interfaces support provisioning and deprovisioning users and groups to an external user store. The IdentityStoreProvisionerWithFiltering interface supports list/query and filtering; the IdentityStoreProvisioner interface does not. For more information about list/query and filtering, see 3.2.2. List/Query Resources and 3.2.2.1. Filtering in the SCIM specification.



Note:

The IdentityStoreUserProvisioner interface is deprecated. Developers should implement either the IdentityStoreProvisionerWithFiltering or IdentityStoreProvisioner interfaces.

IdentityStoreProvisionerWithFiltering interface implementation

Implement the IdentityStoreProvisionerWithFiltering interface to provision and deprovision users and groups to an external user store with list/query and filtering support.



Note:

If you do not need to support list/query and filtering, you can implement the IdentityStoreProvisioner interface instead.

Implementing this interface requires the following Java packages:

- com.pingidentity.sdk.provision
- com.pingidentity.sdk.provision.exception
- com.pingidentity.sdk.provision.users.request
- com.pingidentity.sdk.provision.users.response
- com.pingidentity.sdk.provision.groups.response
- com.pingidentity.sdk.provision.groups.reguest



Group support is optional (see Check for group provisioning support).

For each identity store provisioner implementation, in addition to the methods described under Shared plugin interfaces on page 1025, you must implement the following:

- Create user
- Read user
- Read users (not applicable to the IdentityStoreProvisioner interface)
- Update user
- Delete user
- Check for group provisioning support
- Create group
- Read group
- Read groups (not applicable to the IdentityStoreProvisioner interface)
- Update group
- Delete group

Create user

UserResponseContext createUser(CreateUserRequestContext createRequestCtx) throws IdentityStoreException

PingFederate invokes the createUser() method of your identity store provisioner in response to createuser requests made to PingFederate services, such as inbound provisioning. This method creates the user in the user store managed by the identity store provisioner.

The CreateUserRequestContext contains all information needed to fulfill the request. If the user is successfully provisioned, the method returns a UserResponseContext containing the user attributes used to provision the user. The method throws an IdentityStoreException if an error occurred during the creation process. See the com.pingidentity.sdk.provision.exception package for exceptions that can be thrown.

Read user

UserResponseContext readUser(ReadUserRequestContext readRequestCtx) throws IdentityStoreException

PingFederate invokes the readUser() method of your identity store provisioner in response to get-user requests made to PingFederate services, such as inbound provisioning. This method retrieves user data from the user store managed by the identity store provisioner.

Read users

UsersResponseContext readUsers(ReadUsersRequestContext readRequestCtx) throws IdentityStoreException

PingFederate invokes the readUsers () method of your identity store provisioner in response to list/query requests for user attributes made to PingFederate services, such as inbound provisioning. This method retrieves user data from the user store managed by the identity store provisioner.



Note:

The readUsers method applies only to the IdentityStoreProvisionerWithFiltering interface; it does not apply to the IdentityStoreProvisioner interface.

The ReadUsersRequestContext contains all information needed to fulfill the request. If the user data is successfully retrieved, the method returns a UsersResponseContext containing the user attributes satisfying the filter. If an error occurred during the retrieval process, the method returns an IdentityStoreException . See the com.pingidentity.sdk.provision.exception package for exceptions that can be thrown.

Update user

UserResponseContext updateUser(UpdateUserRequestContext updateRequestCtx) throws IdentityStoreException

PingFederate invokes the updateUser() method of your identity store provisioner in response to updateuser requests made to PingFederate services, such as inbound provisioning. This method updates the user in the user store managed by the identity store provisioner.

The UpdateUserRequestContext contains all information needed to fulfill the request. If the user data is successfully updated, the method returns a UserResponseContext containing the user's updated attributes. The method throws an IdentityStoreException if an error occurred during the update process. See the com.pingidentity.sdk.provision.exception package for exceptions that can be thrown.

Delete user

void deleteUser(DeleteUserRequestContext deleteRequestCtx) throws IdentityStoreException

PingFederate invokes the deleteUser() method of your identity store provisioner in response to deleteuser requests made to PingFederate services, such as inbound provisioning. This method deprovisions the user in the user store managed by the identity store provisioner.

The DeleteUserRequestContext contains all information needed to fulfill the request. The method throws an IdentityStoreException if an error occurred during the deprovision process. See the com.pingidentity.sdk.provision.exception package for exceptions that can be thrown.



Note:

The plugin implementation can choose not to permanently delete the resource, but must return a <code>NotFoundException</code> for all <code>readUser()</code>, <code>updateUser()</code>, and <code>deleteUser()</code> operations associated with the previously deleted ID. In addition, the plugin must not consider the deleted user in conflict calculation. For example, a <code>createUser()</code> request for a user with a previously deleted ID should not throw a <code>ConflictException</code>.

Check for group provisioning support

boolean isGroupProvisioningSupported()
throws IdentityStoreException

Implement the isGroupProvisioningSupported() method to return true if group provisioning is supported by your identity store provisioner or false otherwise. The method throws an IdentityStoreException if an error occurred during the query process. See com.pingidentity.sdk.provision.exception package for exceptions that can be thrown.

Create group

GroupResponseContext createGroup(CreateGroupRequestContext
 createRequestCtx)
throws IdentityStoreException

PingFederate invokes the <code>createGroup()</code> method of your identity store provisioner in response to create-group requests made to PingFederate services, such as inbound provisioning. This method creates the group in the user store managed by the identity store provisioner if the <code>isGroupProvisioningSupported()</code> method returns true; otherwise, it should throw <code>NotImplementedException</code>.

The CreateGroupRequestContext contains all information needed to fulfill the request, such as group attributes. If the group is successfully provisioned, the method returns a GroupResponseContext containing the group attributes used to provision the group. The method throws an IdentityStoreException if an error occurred during the creation process. See the com.pingidentity.sdk.provision.exception package for exceptions that can be thrown.

Read group

 $\label{lem:context} \mbox{ FroupResponseContext readGroup (ReadGroupRequestContext readRequestCtx)} \\ \mbox{ throws IdentityStoreException}$

PingFederate invokes the <code>readGroup()</code> method of your identity store provisioner in response to get-group requests made to PingFederate services, such as inbound provisioning. This method retrieves group data from the user store managed by the identity store provisioner if the <code>isGroupProvisioningSupported()</code> returns true; otherwise, it should throw <code>NotImplementedException</code>.

The <code>ReadGroupRequestContext</code> contains all information needed to fulfill the request, such as group ID. If the group data is successfully retrieved, the method returns a <code>GroupResponseContext</code> containing the group attributes. The method throws an <code>IdentityStoreException</code> if an error occurred during the retrieval process. See the <code>com.pingidentity.sdk.provision.exception</code> package for exceptions that can be thrown.

Read groups

 $\label{lem:groupsResponseContext} \ \ readGroups (ReadGroupsRequestContext \ \ readRequestCtx) \\ throws \ \ IdentityStoreException$

PingFederate invokes the readGroups () method of your identity store provisioner in response to list/query requests for group attributes made to PingFederate services, such as inbound provisioning. This method retrieves group data from the user store managed by the identity store provisioner if the isGroupProvisioningSupported() returns true; otherwise, it should throw NotImplementedException.



Note:

The readGroups method applies only to the IdentityStoreProvisionerWithFiltering interface; it does not apply to the IdentityStoreProvisioner interface.

The ReadGroupsRequestContext will contain all information needed to fulfill the request (for example, a filter). If the group data was successfully retrieved, a GroupsResponseContext should be returned and contain the group attributes for the groups. An IdentityStoreException should be thrown if an error occurred during the retrieval process. See com.pingidentity.sdk.provision.exception package for different exceptions that can be thrown.

Update group

GroupResponseContext updateGroup(UpdateGroupRequestContext updateRequestCtx) throws IdentityStoreException

PingFederate invokes the updateGroup() method of your identity store provisioner in response to update-group requests made to PingFederate services, such as inbound provisioning. This method updates the group in the user store managed by the identity store provisioner if the isGroupProvisioningSupported() method returns true; otherwise, it should throw NotImplementedException.

The UpdateGroupRequestContext contains all information needed to fulfill the request, such as group attributes. If the group data is successfully updated, the method returns a GroupResponseContext containing the group's updated attributes. The method throws an IdentityStoreException if an error occurred during the update process. See the com.pingidentity.sdk.provision.exception package for exceptions that can be thrown.

Delete group

void deleteGroup(DeleteGroupRequestContext deleteRequestCtx) throws IdentityStoreException

PingFederate invokes the deleteGroup() method of your identity store provisioner in response to delete-group requests made to PingFederate services, such as inbound provisioning. This method deprovisions the group in the user store managed by the identity store provisioner if the isGroupProvisioningSupported() returns true; otherwise, it should throw NotImplementedException.

The DeleteGroupRequestContext contains all information needed to fulfill the request, such as a group ID. The method throws an IdentityStoreException if an error occurred during the deprovisioning process. See the com.pingidentity.sdk.provision.exception package for exceptions that can be thrown.

IdentityStoreUserProvisioner interface implementation

The IdentityStoreUserProvisioner interface is deprecated, but you can still implement it to provision and deprovision users to an external user store.

The IdentityStoreUserProvisioner interface is deprecated. Developers can implement it to provision and deprovision users, but they should implement either the IdentityStoreProvisionerWithFiltering or IdentityStoreProvisioner interface.



The IdentityStoreUserProvisioner interface does not provision or deprovision groups. For group support, see IdentityStoreProvisionerWithFiltering interface implementation on page 1034.

The following Java packages are required for implementing the interface:

- com.pingidentity.sdk.provision
- com.pingidentity.sdk.provision.exception
- com.pingidentity.sdk.provision.users.request
- com.pingidentity.sdk.provision.users.response

For each identity store provisioner implementation, in addition to the methods described under Shared plugin interfaces on page 1025, you must implement the following methods:

- Create user
- Read user
- Update user
- Delete user

Create user

UserResponseContext createUser(CreateUserRequestContext createRequestCtx) throws IdentityStoreException

PingFederate invokes the createUser() method of your identity store provisioner in response to createuser requests made to PingFederate services, such as inbound provisioning. This method creates the user in the user store managed by the identity store provisioner.

The CreateUserRequestContext contains all information needed to fulfill the request. If the user is successfully provisioned, the method returns a UserResponseContext containing the user attributes used to provision the user. The method throws an IdentityStoreException if an error occurred during the creation process. See the com.pingidentity.sdk.provision.exception package for exceptions that can be thrown.

Read user

UserResponseContext readUser(ReadUserRequestContext readRequestCtx) throws IdentityStoreException

PingFederate invokes the readUser() method of your identity store provisioner in response to get-user requests made to PingFederate services, such as inbound provisioning. This method retrieves user data from the user store managed by the identity store provisioner.

The ReadUserRequestContext contains all information needed to fulfill the request. If the user data is successfully retrieved, the method returns a UserResponseContext containing the user attributes for the user. The method throws an IdentityStoreException if an error occurred during the retrieval process. See the com.pingidentity.sdk.provision.exception package for exceptions that can be thrown.

Update user

UserResponseContext updateUser(UpdateUserRequestContext updateRequestCtx) throws IdentityStoreException

PingFederate invokes the updateUser() method of your identity store provisioner in response to updateuser requests made to PingFederate services, such as inbound provisioning. This method updates the user in the user store managed by the identity store provisioner.

The UpdateUserRequestContext contains all information needed to fulfill the request. If the user data is successfully updated, the method returns a UserResponseContext containing the user's updated attributes. The method throws an IdentityStoreException if an error occurred during the update process. See the com.pingidentity.sdk.provision.exception package for exceptions that can be thrown.

Delete user

```
void deleteUser(DeleteUserRequestContext deleteRequestCtx)
throws IdentityStoreException
```

PingFederate invokes the deleteUser() method of your identity store provisioner in response to deleteuser requests made to PingFederate services, such as inbound provisioning. This method deprovisions the user in the user store managed by the identity store provisioner.

The DeleteUserRequestContext contains all information needed to fulfill the request. The method throws an IdentityStoreException if an error occurred during the deprovision process. See the com.pingidentity.sdk.provision.exception package for exceptions that can be thrown.



The plugin implementation can choose not to permanently delete the resource, but must return a NotFoundException for all readUser(), updateUser(), and deleteUser() operations associated with the previously deleted ID. In addition, the plugin must not consider the deleted user in conflict calculation. For example, a createUser() request for a user with a previously deleted ID should not throw a ConflictException.

Developing notification publishers

To develop a notification publisher, implement the NotificationPublisherPlugin interface.

PingFederate administrators can define and manage notification publishers, as described in *Managing* notification publisher instances on page 926. If those features do not meet your needs, you can develop a custom notification publisher using the PingFederate NotificationPublisherPlugin interface and the following Java packages:

- com.pingidentity.sdk
- org.sourceid.saml20.adapter.conf
- org.sourceid.saml20.adapter.gui

The NotificationPublisherPlugin interface, which extends the Plugin interface, defines the publishNotification() method.

For each implementation, define the publishNotification () method, in addition to the methods described in Shared plugin interfaces on page 1025. PingFederate invokes the publishNotification() method when publishing a notification. For example, you can configure PingFederate so that an account password change invokes the method.

```
PublishResult publishNotification(String eventType,
 Map<String, String> data,
```

Map<String, String> configuration)

The method returns the PublishResult, the status of the notification that the plugin instance sent.

For more information about the NotificationPublisherPlugin interface, see the SDK Javadocs. You can also see a sample implementation in pingfederate/sdk/plugin-src.

Building and deploying with Ant

Use the Apache Ant build script to clean, build, package, and deploy projects within the PingFederate Java SDK.

About this task

The PingFederate Java SDK comes with an Apache Ant build script that simplifies building and deploying your projects.

Steps

1. Edit the build.local.properties file and set the target-plugin.name property to the name of your project subdirectory (see *Directory structure*).



Note:

You can develop source code for multiple projects simultaneously, but you can build and deploy only one at a time. Change the value of the target-plugin.name property as needed to build and deploy other projects.

- 2. If your project depends on any third-party .jar files, place them into your project's lib directory. If the directory does not exist, create a new directory called 11b directly under your project's directory. For example, pingfederate/sdk/plugin-src/<subproject-name>/lib.
- 3. On the command line in the sdk directory, use ant to clean, build, and package or to build, package, and deploy your project.

Option	Description
Clean the project	ant clean-plugin
Compile the project	ant compile-plugin
Compile the project and create a .jar file	ant jar-plugin



Note:

The SDK creates a deployment descriptor in the PF INF directory and places it in a .jar file. The descriptor tells PingFederate what plugin implementations are contained in the file, and the compiled class files and the deployment descriptor are placed in the pingfederate/sdk/pluginsrc/<subproject-name>/build/classes directory. The pf.plugins.<subprojectname>.jar file is placed in the pingfederate/sdk/plugin-src/<subproject-name>/build/ jar directory.

To compile the project, create a .jar file, and deploy the project to PingFederate, enter:

```
ant deploy-plugin
```

This build target performs the steps described above and deploys any .jar files found in the lib directory of your subproject.



Note:

To deploy your plugin manually to an installation of the PingFederate server, copy the .jar file and any third-party.jar files into the /server/default/deploy/ directory of that PingFederate installation.

4. Restart the PingFederate server.

Building and deploying manually

Use a build utility to add directories, create deployment descriptors, and create a .jarfile to build and deploy your plugins with PingFederate.

Before you begin

To compile your project, you must have the following directories on your classpath:

- <pf install>/pingfederate/server/default/lib
- <pf install>/pingfederate/lib
- <pf install>/pingfederate/sdk/lib
- <pf install>/pingfederate/sdk/plugin-src/<subproject-name>/lib

About this task

To build your project with another build utility, you must create the deployment descriptors for each of your plugins. The deployment descriptor files allow PingFederate to discover your plugins. Once this is complete, use the build tool to create a .jar file and deploy it within the appropriate directory.

Steps

- 1. Add a new directory called PF-INF into your project. This directory must be at the root of your .jar file, similar to META-INF.
- 2. In PF-INF, add an appropriate text file for each type of plugin you created:

Plugin type	File name
IdP Adapter	idp-authn-adapters
SP Adapter	sp-authn-adapters
Custom Data Source	custom-drivers
Token Processor	token-processors
Token Generator	token-generators
Authentication Selector	authentication-selectors
Password Credential Validator	password-credential-validators

Plugin type	File name
Identity Store Provisioner	identity-store-provisioners
CIBA Authenticator	oob-auth-plugins
Notification Publisher	notification-sender

- 3. In each text file added, specify the fully-qualified class name of each plugin that implements the corresponding plugin interface. Place each class name on a separate line.
- 4. To create a .jar, archive the compiled class files along with the deployment descriptors using your build tool. The deployment descriptors must be in the PF-INF directory, located at the root of the .jar file.
- 5. To deploy your plugin, copy the .jar file and any third-party .jar files into the <pf_install>/ pingfederate/server/default/deploy directory of the PingFederate installation.

Log messages

Use a typical logging pattern based on the Apache Commons logging framework to log messages from your adapter, token translator, or custom data source driver.

The service provider (SP) adapter contained in the directory sdk/adapters-src/sp-adapter-example shows how to use a logger for your adapter.

Developer's Reference Guide

This document describes the PingFederate endpoints and APIs.

OAuth 2.0 endpoints

When developing OAuth-capable applications, developers must follow the OAuth 2.0 Authorization Framework and OpenID Connect specifications if applicable.

OAuth-capable applications must send requests to various OAuth endpoints to obtain authorization grants, access tokens, refresh tokens, and ID tokens if applicable. Additional endpoints exist for other purposes, including for clients to validate access and refresh tokens, for developers to submit client registrations using the OAuth 2.0 Dynamic Client Registration protocol, and for clients to retrieve OpenID Connect metadata.

Each endpoint extends from the runtime server at the base URL. If you configure virtual host names, the endpoints are also accessible at those locations.

Example

Example

If the base URL is https://www.example.com:9031 and the configured virtual host names are www.example.org and www.example.info, the authorization and token endpoints are accessible at the following locations:

Authorization endpoint /as/authorization.oauth2

- https://www.example.com:9031/as/authorization.oauth2
- https://www.example.org:9031/as/authorization.oauth2

https://www.example.info:9031/as/authorization.oauth2

Token endpoint /as/token.oauth2

- https://www.example.com:9031/as/token.oauth2
- https://www.example.org:9031/as/token.oauth2
- https://www.example.info:9031/as/token.oauth2

Authorization endpoint

The OAuth authorization server (AS) uses the authorization endpoint to interact directly with resource owners, authenticate them, and obtain their authorizations.

The OAuth 2.0 Authorization Framework defines the authorization endpoint. Typically, an OAuth client makes an authorization request by directing a resource owner through an HTTP user-agent to the authorization endpoint. After the OAuth AS completes its interaction with the resource owner, the OAuth AS redirects the resource owner's user-agent back to the client's redirect URI with the response to the authorization request.



This endpoint can be used in an OAuth Scope Authentication Selector configuration, which can affect the behavior of the endpoint. For example, the idp parameter might be enforced or overridden by policy determined by an instance of the OAuth Scope Authentication Selector.

This endpoint accepts the HTTP GET and POST methods.

Endpoint: /as/authorization.oauth2

When transmitting through the HTTP POST method, the required Content-Type value is application/x-www-form-urlencoded. The following table describes parameters for this endpoint.

Parameter	Description
client_id	The client identifier.
(Required)	
response_mode	When set to <code>form_post</code> , the authorization response is returned to the client in an auto-POST form, in accordance with the OAuth 2.0 Form Post Response Mode specification.
	When set to pi.flow, if PingFederate is configured to let OAuth clients initiate authorization directly through the authentication API without needing to handle HTTP redirections, then the client receives an authorization response with a flow ID and other information it needs to continue authentication. For more information, see <i>Mobile application authentication through REST APIs</i> on page 1159.
response_type (Required)	A value of code results in the Authorization Code grant type while a value of token implies the Implicit grant type. Additionally, a value of id_token can be requested by implicit clients.
	To initiate a Hybrid Flow, multiple response_type values can be specified by space-separating them. When using the Hybrid Flow, some tokens are returned from the Authorization Endpoint and others are returned from the Token Endpoint. For information about multiple-valued response type combinations, see the description of the restrictedResponseTypes parameter in OAuth Client Management Service on page 1093.

Parameter

Description

code challenge

To reduce the risk of authorization code interception attack, supply a one-time string value to associate the authorization request with the token request. For more information, see Proof Key for Code Exchange (PKCE) by OAuth Public Clients.

Applicable only when response type parameter value is code. Mandatory if the client is required to do so. For more information, see Require Proof Key for Code Exchange (PKCE) in Configuring an OAuth client on page 515.



Note:

If used, the OAuth client must submit the corresponding code verifier when using the authorization code to obtain an access token. For more information, see code verifier in OAuth grant type parameters on page 1060.

code challenge methoapplicable only when the response type parameter value is code and a code challenge parameter value is provided.

> This parameter indicates the transformation method used to derive the code challenge parameter value from that of the code verifier parameter. PingFederate OAuth AS supports two transformation methods:

- plain, which indicates the code challenge parameter value is that of the code verifier parameter.
- S256, which indicates the code challenge parameter is derived from the code verifier parameter value as follows

code challenge=Base64Urlencode (SHA256 (ASCII (code_verifier))), where:

- ASCII (code verifier) denotes the octets of the ASCII representation of the code verifier value.
- SHA256 (octets) denotes the SHA 256-bit hash of the octets.
- Base64Url-encode (octets) denotes the base64url encoding of octets; the output is URL-safe.



Note:

For detailed information about the transformation method, see *Proof Key* for Code Exchange (PKCE) by OAuth Public Clients.

The code challenge method parameter value is case-sensitive. An error message is returned to the clients for any other values.

Omitting the code challenge method parameter has the same effect as providing the code challenge method parameter with a value of plain.

Parameter Description The URI to which PingFederate redirects the resource owner's user-agent redirect uri after an authorization is obtained. For OpenID Connect protocol compliance, clients that use the authorization code or implicit grant type must include this parameter in their authorization requests. It is also the default behavior in new PingFederate installations starting with version 9.1.4. For upgraded installations, this requirement remains true for clients that have been configured with more than one redirection URIs. For clients that have been configured with only one redirection URI, this requirement is waived to minimize the impact that it might impose on customers upgrading to version 9.1.4 or a subsequent release. As needed, it can be enabled at a later time. 1 Note: If this parameter is used, the same parameter and value must also be used in subsequent token requests. For more information, see OAuth grant type parameters on page 1060. Specifies the end-user's preferred languages for claims being returned in a claims locales space-separated list, ordered by preference. The values must conform to the IETF BCP 47 guidelines. Tip: You can map the claims locales value into the persistent grants (and therefore the access tokens, the ID tokens, or both) from an IdP adapter or an IdP connection by selecting Context under Source and Requested Claims Locales under Value in the Contract Fulfillment tab in the IdP Adapter Mapping configuration or the OAuth Attribute Mapping configuration in an IdP connection. login_hint Provides a hint to the PingFederate AS about the end user. For example, when an OAuth client includes a login hint in its authorization request and the authentication source is an HTML Form Adapter instance, the username field in the login form is pre-populated with the login hint parameter value. Sets an allowable elapsed time in seconds since the end user last max age authenticated. If the elapsed time exceeds the value of max age, PingFederate prompts the end user for authentication. **Tip:**

The HTML Form Adapter supports the max age parameter by tracking the

authentication time for each user.

Parameter Description

request

A single, self-contained parameter; a signed JWT whose claims represent the request parameters of the authorization request. The OpenID Connect specification calls this JWT a request object.

The request parameter is required if a client is configured to transmit request parameters in signed request objects. When PingFederate receives an authorization request, it verifies the digital signature of the signed request object based on the key obtained from the pre-configured JWKS URL or JWKS, and the selected request object signing algorithms. If the signature does not pass the verification process, the request fails.

The request parameter is optional if a client is not configured to transmit request parameters in signed request objects but is configured with a JWKS URL or an actual JWKS. This flexibility allows the client to transmit request parameters in signed request objects for some requests and without the use of signed request objects for some other transactions. When PingFederate receives an authorization request with a signed request object, it verifies the digital signature of the signed request object based on the key obtained from the pre-configured JWKS URL or JWKS, and the selected request object signing algorithms. If the signature does not pass the verification process, the request fails.

If a client is not configured to transmit request parameters in signed request objects and is not configured with a JWKs URL or an actual JWKs, PingFederate ignores the request parameter. When PingFederate receives an authorization request with a signed request object, it processes the authorization request and disregards the signed request object. As needed, develop a custom IdP adapter using the PingFederate SDK to extract the request parameter and its value from the HTTP request for further processing.



Note:

If a client includes in an authorization request a request parameter, other than client id and response type, as a parameter outside of the signed request object and a claim inside of the signed request object, PingFederate always uses the claim value found inside the signed request object to process the request further.

For the client id and response type request parameters, the values outside of the signed request object must match the claim values inside of the signed request object. If the values do not match, PingFederate returns an error message to the client.

If a request parameter is found only outside of the signed request object, PingFederate ignores the request parameter and returns no error message.



Per OAuth and OpenID Connect specifications, a client must always include in an authorization request the client id, response type, and scope request parameters outside of the signed request object.

For client configuration information, see the Require Signed Request setting in Configuring an OAuth client on page 515. For more information about request objects, see OpenID Connect specifications.

Parameter	Description
scope	Expresses the scope of the access request as a list of space-separated, case-sensitive strings. For detailed information about scopes, see <i>Scopes and scope management</i> on page 486.
state	An opaque value used by the client to maintain state between the request and callback. If included, the AS returns this parameter and the given value when redirecting the user agent back to the client.
ui_locales	Specifies the end-user's preferred languages for OAuth user interactions in a space-separated list, ordered by preference. The values must conform to the <i>IETF BCP 47</i> guidelines.
idp Of PartnerIdpId	A PingFederate OAuth AS parameter indicating the entity ID or the connection ID of the IdP with whom to initiate Browser SSO for user authentication.
pfidpadapterid Or IdpAdapterId	A PingFederate OAuth AS parameter indicating the IdP adapter instance ID of the adapter to use for user authentication.
	Note: This parameter might be overridden by policy based on authentication policies. For example, an OAuth Scope Authentication Selector instance could enforce the use of a given adapter instance based on client-requested scopes.

If more than one source of authentication is configured in the system and no pfidpadapterid or idp parameter is provided, PingFederate provides users with an intermediate page asking them to choose among the available sources of authentication. The authentication results in a set of user attributes that must be mapped into the USER_KEY attribute for persistent grant storage and the USER_NAME attribute that displays on the user authorization page.

OpenID Connect parameters

The following table describes OpenID Connect parameters for this endpoint.

Parameter	Description
acr_values	Specifies the Authentication Context Class Reference (acr) values for the AS to use when processing an Authentication Request. Express the values as a space-separated string, and list them in order of preference.
id_token_hint	Includes an ID token as a hint to the PingFederate AS about the end user. If the authenticated user does not match the information stored in the ID token, the PingFederate AS rejects the authorization request and returns an error message.
nonce	Specifies a string value used to associate a client session with an ID token and to reduce replay attacks. The value passes through unmodified from an authorization request to the ID token.
prompt	Specifies whether the AS prompts the end user for reauthentication and consent. Expressed as a list of space-separated, case-sensitive ASCII string values. If included, the client can use this parameter to verify that the end user is still present for the current session or to bring attention to the request.
	PingFederate supports values of none, login, and consent.

OAuth access token management parameters

PingFederate supports multiple access token management (ATM) instances. Clients can specify an ATM instance by providing the ATM ID (access_token_manager_id) or a resource URI (aud) in their requests to the PingFederate OAuth AS.

Parameter	Description
access_token_manage	arThelaccess_token_manager_id value is the instance ID of the desired ATM instance. When specified, PingFederate uses the desired ATM instance for the request if it is eligible; otherwise it aborts the request.
	Note: When the access_token_manager_id parameter is specified, PingFederate ignores the aud parameter.
aud	The aud is the resource URI the client wants to access. The provided value is matched against resource URIs configured in access token management instances. When a match is found, PingFederate uses the corresponding access token management instance for the request if it is eligible; otherwise it aborts the request.

A match can be an exact match or a partial match where the provided URI has the same scheme and authority parts and a more specific path contained within the path of the pre-configured resource URI. PingFederate takes an exact match over a partial match. If there are multiple partial matches, PingFederate takes the partial match where the provided URI matches more specifically against the preconfigured resource URI.

Example

A partial match

A resource URI of https://app.example.local is a partial match for the following provided URIs:

- https://app.example.local/file1.ext
- https://app.example.local/path/file2.ext
- https://app.example.local/path/more

Example

An exact match is a better match than a partial match

Access Token Management instances	Resource URIs (configured)
ATM1	https://localhost:9031/app1
	https://localhost:9031/app2/data
	https://app.example.local
ATM2	https://localhost:9031/app1/data
	https://localhost:9031/app2/data/get

https://localhost:9031/app1 (a resource URI pre-configured for ATM1) is a partial match for https://localhost:9031/app1/data (the provided URI). However, PingFederate chooses ATM2 because https://localhost:9031/app1/data (a resource URI pre-configured for ATM2) is an exact match against the provided URI.

Example

A more specific partial match is a better match

Both https://localhost:9031/app2/data (a resource URI for ATM1) and https:// localhost: 9031/app2/data/get (a resource URI for ATM2) are partial matches for https:// localhost:9031/app2/data/get/sample (the provided URI). However, PingFederate chooses ATM2 because https://localhost:9031/app2/data/get matches more specifically against the provided URI.

Client-initiated backchannel authentication endpoint

A CIBA-capable client uses this endpoint to initiate a backchannel, out-of-band flow to authenticate the resource owners and obtain their authorizations.

The OpenID Connect Client Initiated Backchannel Authentication Flow defines the client-initiated backchannel authentication (CIBA) endpoint.



This endpoint accepts only the HTTP POST method.

Endpoint: /as/bc-auth.ciba

The following table describes parameters for this endpoint. The required Content-Type value is application/x-www-form-urlencoded.

Parameter	Description
client_id	The client identifier.
(Required)	Important:
	When sending request parameters of an authentication request with a signed request object, the client must include the client_id parameter and its value inside and outside of the request parameter value. Both client_id parameter values must match.
scope (Required)	The scope of the access request. Expressed as a list of space-separated, case-sensitive strings.
	Scope values are globally defined on the System # OAuth Settings # Scope Management window. You can constrain scopes on a client-to-client basis.
	This parameter must include the openid scope value.
client_notification	Abberer token provided by the client that PingFederate must include when sending a ping callback message to the client's notification endpoint. This usage must conform to the syntax for bearer credentials as defined in section 2.1 of RFC 6750.
	If the client is configured to use the poll delivery method, this parameter is required.

Parameter	Description
id_token_hint, login_hint_token, or login_hint	Per the CIBA specification, the client must include one and only one hint for the OpenID Provider to identify the user. The valid hint parameters are id_token_hint, login_hint, and login_hint_token.
	id_token_hint
	Use this parameter to include an ID token as a hint for PingFederate to identify the user. This ID token must be unencrypted. It must be a signed ID token.
	login_hint_token
	Use this parameter to include a JSON web token (JWT) as a hint for PingFederate to identify the user. The attributes of this token can vary from one use case to another. For more information how PingFederate uses the login hint token, see <i>Configuring identity hint contract</i> on page 583.
	login_hint
	Use this parameter to provide a hint to PingFederate to identify the user. The value can contain an email address, phone number, account number, subject identifier, username, or any attribute that both sides agreed upon.
user_code	A secret code that is known only to the user and verifiable by PingFederate through the use of a Password Credential Validator instance. The purpose of this code is to authorize the transmission of an authentication request to the user's authentication device.
	If the client record is configured to support user code and associated with a user code-enabled CIBA request policy, this parameter is required.
binding_message	An alphanumeric message intended to be made available on both the authentication device and the consumption device. The user can tie them together and decide whether to grant the authorization.
	When provided, the length of the message must range from 1 - 20 characters.
requested_expiry	The requested expiration time of the request in seconds since the generation of the authentication request acknowledgment.
	Note:
	PingFederate honors the requested expiration time only if the value is shorter than that of the Transaction Lifetime field found in the associated CIBA request policy.

Sample authentication request

```
POST /as/bc-auth.ciba HTTP/1.1
Content-Type: application/x-www-form-urlencoded
Host: www.example.com
client_id=myCibaApp&scope=openid&login_hint=joe@example.com
```

with a signed request object will always fail.

configured with a JWKS URL or an actual JWKS, an authentication request

Sample authentication request acknowledgements

• 200 - Success

```
HTTP/1.1 200 OK
...
{
    "auth_req_id":
    "HQnCgSeUzWNORZEv8n3E8wIip9L3iwBdJAAect04BqdpEsFBGqfxRvoa_Q",
        "interval": 3,
        "expires_in": 120
}
```

400 - Bad Request

```
HTTP/1.1 400 Bad Request
   "error description": "CIBA authentication requests MUST contain
the openid scope value.",
   "error": "invalid scope"
HTTP/1.1 400 Bad Request
    "error description": "Authentication request parameters (such
as binding message) MUST NOT be present outside of the JWT when a
signed authentication request is used.",
    "error": "invalid request"
HTTP/1.1 400 Bad Request
    "error description": "Exactly one (not more, not less) of
the hint parameters (i.e. 'login hint token', 'id token hint' or
 'login hint') must be provided.",
    "error": "invalid request"
HTTP/1.1 400 Bad Request
   "error description": "User could not be sufficiently identified
to initiate out-of-band auth",
   "error": "unknown user id"
}
HTTP/1.1 400 Bad Request
{ "error": "invalid user code" }
HTTP/1.1 400 Bad Request
{ "error": "missing user code" }
HTTP/1.1 400 Bad Request
    "error description": "Client is not configured to support user
code but a user code was sent in the request.",
    "error": "invalid request"
HTTP/1.1 400 Bad Request
    "error description": "Policy is set to require a token for the
user hint but login hint was sent.",
```

```
"error": "invalid_request"
}
```

• 401 - Unauthorized

```
HTTP/1.1 401 Unauthorized
...
{
    "error_description": "Invalid client or client credentials.",
    "error": "invalid_client"
}
```

• 500 - Server Error

```
HTTP/1.1 500 Server Error
...
{
    "error_description": "Client is configured to support user code
but server policy doesn't have a PCV configured to do the user code
checking",
    "error": "server_error"
}
```

For more information about error responses, see section 13. Authentication Error Response in the specification.

OAuth client identification and authentication

The authentication requirement of this endpoint depends on the client authentication method configured for the clients.

Authentication method	Parameters
Client secret	Clients can present their client identifier and client secret using the HTTP Basic authentication scheme, where the client identifier is the username, and the client secret is the password.
	Clients can provide credentials using the request parameters client_id and client_secret.
	Important:
	This is a sensitive parameter. To avoid recording it in web server logs, only pass in this parameter with the HTTP POST method in the message body, or through the HTTP Basic authentication scheme.
Client certificate	Clients must present their client certificate for mutual TLS authentication. The issuer and the subject distinguished name (DN) of the client certificate must match values configured for the clients.

Authentication method	Parameters
Private key JWT	Clients must include request parameters client_assertion_type and client_assertion in the message body of their requests.
	client_assertion_type
	The value describes the format of the assertion as defined by the authorization server. For the private_key_jwt client authentication method, the value is urn:ietf:params:oauth:client-assertion-type:jwt-bearer.
	client_assertion
	The value is the authentication token.
	Example
	<pre>client_assertion_type= urn%3Aietf%3Aparams%3Aoauth% 3Aclient-assertion-type%3Ajwt-bearer& client_assertion= eyJhbGciOiJSUzI1NiIsLbSWi1YO-TILOd4L7ZCg&</pre>
	Note:
	For readability, line breaks are inserted and the authentication token is truncated.
	For more information about the private_key_jwt client authentication method, see Client Authentication and Using Assertions for Client Authentication.
None	Clients must pass in the client_id parameter in a query string or the message body to identify themselves.

OAuth access token management parameters

PingFederate supports multiple access token management (ATM) instances. Clients can specify an ATM instance by providing the ATM ID (access_token_manager_id) or a resource URI (aud) in their requests to the PingFederate OAuth AS.

Parameter	Description
access_token_manage	rThetaccess_token_manager_id value is the instance ID of the desired ATM instance. When specified, PingFederate uses the desired ATM instance for the request if it is eligible; otherwise it aborts the request.
	Note:
	When the access_token_manager_id parameter is specified, PingFederate ignores the aud parameter.

Parameter	Description
aud	The aud is the resource URI the client wants to access. The provided value is matched against resource URIs configured in access token management instances. When a match is found, PingFederate uses the corresponding access token management instance for the request if it is eligible; otherwise it aborts the request.

A match can be an exact match or a partial match where the provided URI has the same scheme and authority parts and a more specific path contained within the path of the pre-configured resource URI. PingFederate takes an exact match over a partial match. If there are multiple partial matches, PingFederate takes the partial match where the provided URI matches more specifically against the pre-configured resource URI.

Example

A partial match

A resource URI of https://app.example.local is a partial match for the following provided URIs:

- https://app.example.local/file1.ext
- https://app.example.local/path/file2.ext
- https://app.example.local/path/more

Example

An exact match is a better match than a partial match

Access Token Management instances	Resource URIs (configured)
ATM1	https://localhost:9031/app1
	https://localhost:9031/app2/data
	https://app.example.local
ATM2	https://localhost:9031/app1/data
	https://localhost:9031/app2/data/get

https://localhost:9031/app1 (a resource URI pre-configured for ATM1) is a partial match for https://localhost:9031/app1/data (the provided URI). However, PingFederate chooses ATM2 because https://localhost:9031/app1/data (a resource URI pre-configured for ATM2) is an exact match against the provided URI.

Example

A more specific partial match is a better match

Both https://localhost:9031/app2/data (a resource URI for ATM1) and https://localhost:9031/app2/data/get (a resource URI for ATM2) are partial matches for https://localhost:9031/app2/data/get/sample (the provided URI). However, PingFederate chooses ATM2 because https://localhost:9031/app2/data/get matches more specifically against the provided URI.

Token endpoint

The client presents its authorization grant to the token endpoint to obtain an access token and a refresh token when needed.

The OAuth 2.0 Authorization Framework defines the token endpoint. Because the authorization endpoint directly issues an access token, every authorization grant uses this endpoint except the Implicit grant type.



Note:

This endpoint accepts only the HTTP POST method.

Endpoint: /as/token.oauth2

Parameters vary depending on the grant type. For more information, see OAuth grant type parameters on page 1060. The required Content-Type value is application/x-www-form-urlencoded.

Like other OAuth 2.0 endpoints, the token endpoint is accessible at the base URL and any configured virtual host names.

If the Token Endpoint Base URL field is configured on the Authorization Server Settings window (System # OAuth Settings # Authorization Server Settings) the token endpoint is also accessible at that location.

For example, if the base URL is https://www.example.com:9031 and the Token Endpoint Base URL value is https://www.example.local:9031, the token endpoints are accessible at the following locations:

- https://www.example.com:9031/as/token.oauth2
- https://www.example.local:9031/as/token.oauth2

OAuth client identification and authentication

The authentication requirement of this endpoint depends on the client authentication method configured for the clients.

Authentication method	Parameters
Client secret	Clients can present their client identifier and client secret using the HTTP Basic authentication scheme, where the client identifier is the username, and the client secret is the password.
	Clients can provide credentials using the request parameters client_id and client_secret.
	Important:
	This is a sensitive parameter. To avoid recording it in web server logs, only pass in this parameter with the HTTP POST method in the message body, or through the HTTP Basic authentication scheme.
Client certificate	Clients must present their client certificate for mutual TLS authentication. The issuer and the subject distinguished name (DN) of the client certificate must match values configured for the clients.

Authentication method	Parameters
Private key JWT	Clients must include request parameters client_assertion_type and client_assertion in the message body of their requests.
	client_assertion_type
	The value describes the format of the assertion as defined by the authorization server. For the private_key_jwt client authentication method, the value is urn:ietf:params:oauth:client-assertion-type:jwt-bearer.
	client_assertion
	The value is the authentication token.
	Example
	<pre>client_assertion_type= urn%3Aietf%3Aparams%3Aoauth% 3Aclient-assertion-type%3Ajwt-bearer& client_assertion= eyJhbGciOiJSUzI1NiIsLbSWi1YO-TILOd4L7ZCg&</pre>
	Note:
	For readability, line breaks are inserted and the authentication token is truncated.
	For more information about the private_key_jwt client authentication method, see Client Authentication and Using Assertions for Client Authentication.
None	Clients must pass in the client_id parameter in a query string or the message body to identify themselves.

OAuth access token management parameters

PingFederate supports multiple access token management (ATM) instances. Clients can specify an ATM instance by providing the ATM ID (access_token_manager_id) or a resource URI (aud) in their requests to the PingFederate OAuth AS.

Parameter	Description
access_token_manage	arThetaccess_token_manager_id value is the instance ID of the desired ATM instance. When specified, PingFederate uses the desired ATM instance for the request if it is eligible; otherwise it aborts the request.
	Note:
	When the access_token_manager_id parameter is specified, PingFederate ignores the aud parameter.

Parameter	Description
aud	The aud is the resource URI the client wants to access. The provided value is matched against resource URIs configured in access token management instances. When a match is found, PingFederate uses the corresponding access token management instance for the request if it is eligible; otherwise it aborts the request.

A match can be an exact match or a partial match where the provided URI has the same scheme and authority parts and a more specific path contained within the path of the pre-configured resource URI. PingFederate takes an exact match over a partial match. If there are multiple partial matches, PingFederate takes the partial match where the provided URI matches more specifically against the pre-configured resource URI.

Example

A partial match

A resource URI of https://app.example.local is a partial match for the following provided URIs:

- https://app.example.local/file1.ext
- https://app.example.local/path/file2.ext
- https://app.example.local/path/more

Example

An exact match is a better match than a partial match

Access Token Management instances	Resource URIs (configured)
ATM1	https://localhost:9031/app1
	https://localhost:9031/app2/data
	https://app.example.local
ATM2	https://localhost:9031/app1/data
	https://localhost:9031/app2/data/get

https://localhost:9031/app1 (a resource URI pre-configured for ATM1) is a partial match for https://localhost:9031/app1/data (the provided URI). However, PingFederate chooses ATM2 because https://localhost:9031/app1/data (a resource URI pre-configured for ATM2) is an exact match against the provided URI.

Example

A more specific partial match is a better match

Both https://localhost:9031/app2/data (a resource URI for ATM1) and https://localhost:9031/app2/data/get (a resource URI for ATM2) are partial matches for https://localhost:9031/app2/data/get/sample (the provided URI). However, PingFederate chooses ATM2 because https://localhost:9031/app2/data/get matches more specifically against the provided URI.

The /as/token.oauth2 endpoint accepts other parameters which vary by the grant type presented.

These parameters include OAuth-defined standard parameters and parameters proprietary to PingFederate. The following parameter indicates the grant type of the access token request.

Parameter	Description
grant_type	Indicates the type of grant being presented in exchange for an access token
(Required)	and possibly a refresh token. The value is an extensibility mechanism of the OAuth 2.0 specification. PingFederate supports these values:
	authorization_code
	refresh_token
	• password
	client credentials
	urn:openid:params:grant-type:ciba
	urn:ietf:params:oauth:grant-type:device code
	urn:ietf:params:oauth:grant-type:jwt-bearer
	urn:ietf:params:oauth:grant-type:saml2-bearer
	urn:ietf:params:oauth:grant-type:token-exchange
	urn:pingidentity.com:oauth2:grant_type:validate_bearer
	Note: The following sections define further parameters associated with each grant type.

Authorization code grant type

These parameters apply when the <code>grant_type</code> parameter for <code>/as/token.oauth2</code> is set to authorization_code.

Parameter	Description
code (Required)	The authorization code received from the authorization server during the redirect interaction at the authorization endpoint when the response_type parameter is code.
code_verifier	Required if the authorization request was sent with a code_challenge parameter to reduce the risk of code interception attack.
	If a code_challenge_method parameter value is provided in the request for the authorization code, PingFederate OAuth Authorization Server (AS) validates the code_verifier parameter value against that of the code_challenge value. If the validation returns no errors, PingFederate OAuth AS returns an access token; otherwise it returns an error to the client.
	For more information about the code_challenge parameter, the code_challenge_method parameter, and the support for the <i>Proof Key for Code Exchange (PKCE) by OAuth Public Clients</i> specification (tools.ietf.org/html/rfc7636), see <i>Authorization endpoint</i> on page 1044.

Parameter	Description
redirect_uri	This parameter is required if the redirect_uri parameter was included in the authorization request that resulted in the issuance of the code. For more information, see <i>Authorization endpoint</i> on page 1044). The value here must match the authorization-request value, if applicable.
	The parameter is also required for clients with multiple redirection URIs or one redirection URI that uses wildcards.
	The parameter is optional for clients with only one specific redirection URI.
(Optional) sens	The scope of the access request expressed as a list of space-delimited, case- sensitive strings. The requested scope must be equal to or less than the scope
	originally authorized. If omitted, the scope is treated as equal to the scope originally authorized.
	Scopes can also be constrained on a client-to-client basis. For more information, see <i>Scopes and scope management</i> on page 486.

Refresh token grant type

These parameters apply when the $\tt grant_type$ parameter for /as/token.oauth2 is set to refresh token.

Parameter	Description
refresh_token	The refresh token issued to the client during a previous access token request.
(Required)	Important:
	To avoid recording this parameter in web server logs, only pass it in the message body using the HTTP POST method.
scope	The scope of the access request expressed as a list of space-delimited, case- sensitive strings. The requested scope must be equal to or less than the scope originally granted. If omitted, the scope is treated as equal to the original set.
	Scopes can also be constrained on a client-to-client basis. For more information, see <i>Scopes and scope management</i> on page 486.

Resource owner password credentials grant type

 $\textbf{These parameters apply when the grant_type parameter for } / \texttt{as/token.oauth2} \textbf{ is set to } \texttt{password.}$

Parameter	Description
username	The username, encoded as UTF-8.
(Required)	
password	The password, encoded as UTF-8.
(Required)	Important:
	To avoid recording this parameter in web server logs, only pass it in the message body using the HTTP POST method.

When a token request triggers an invalid grant error because the corresponding LDAP Username Password Credential Validator instance returns an authentication error, PingFederate includes the relevant message in the error response. See the following example.

```
"error": "invalid grant",
 "error description": "We didn't recognize the username or password you
entered. Please try again."
```

The error description varies based on the error condition that the LDAP Username PCV detects. OAuthclient developers can create custom experiences based on the error messages.



These customizable messages are stored in the PingFederate message file, pingfederatemessages.properties, located in the cpf install>/pingfederate/server/default/conf/ language-packs directory.

You can localize these messages by using the PingFederate localization framework for an international audience. For more information, see .

The client id parameter is not required when the Allow unidentified clients to make resource owner password credentials grants check box is selected in the System # OAuth Settings # Authorization Server Settings window.

Client credentials grant type

The following parameters applies when the grant type parameter for /as/token.oauth2 is set to client credentials.

Parameter	Description
scope	The scope of the access request expressed as a list of space-delimited, case-sensitive strings. Scopes can also be constrained on a client-to-client basis. For more information, see <i>Scopes and scope management</i> on page 486.

Client-initiated backchannel authentication (CIBA)

The following parameter applies when the grant type parameter for /as/token.oauth2 is set to urn:openid:params:grant-type:ciba.

Parameter	Description
auth_req_id	The unique identifier to identify the authentication request.

Sample request

```
POST /as/token.oauth2 HTTP/1.1
Content-Type: application/x-www-form-urlencoded
Host: www.example.com
...
grant_type=urn%3Aopenid%3Aparams%3Agrant-type
%3Aciba&client_id=myCibaApp&auth_req_id=yQn...1Vw
```

The auth req id parameter value in the sample is truncated for readability.

Device authorization grant type

The following parameter applies when the grant_type parameter for /as/token.oauth2 is set to urn:ietf:params:oauth:grant-type:device code.

The *OAuth 2.0 Device Authorization Grant* specification defines the process that allows a user to grant authorization to a device using a browser on a second device, such as a smart phone or a computer. For more information, see *Device authorization grant* on page 51.

Parameter	Description
device_code	The device code found in the device authorization response.
(Required)	

JWT Bearer Token grant type

These parameters apply when the grant_type parameter for /as/token.oauth2 is set to urn:ietf:params:oauth:grant-type:jwt-bearer.

Parameter	Description
assertion	A JSON Web Token (JWT), as defined in RFC7523, section 2.1.
(Required)	Important: To avoid recording this parameter in web server logs, only pass it in the message body using the HTTP POST method.
scope	The scope of the access request expressed as a list of space-delimited, case-sensitive strings. Scopes can be constrained on a client-to-client basis. For more information, see <i>Scopes and scope management</i> on page 486.

The client_id parameter is not required when the Allow unidentified clients to request extension grants check box is selected in the System # OAuth Settings # Authorization Server Settings window.

SAML 2.0 Bearer Assertion grant type

The following parameters apply when the grant_type parameter for /as/token.oauth2 is set to urn:ietf:params:oauth:grant-type:sam12-bearer.

Parameter	Description
assertion (Required)	A single SAML 2.0 assertion, which must be encoded using base64url. For more information, see <i>RFC4648</i> , section 5.
	Important:
	To avoid recording this parameter in web server logs, only pass it in the message body using the HTTP POST method.
scope	The scope of the access request expressed as a list of space-delimited, case-sensitive strings. Scopes can be constrained on a client-to-client basis. For more information, see <i>Scopes and scope management</i> on page 486.



The client_id parameter is not required when the Allow unidentified clients to request extension grants check box is selected in the System # OAuth Settings # Authorization Server Settings window.

token exchange grant type

The following parameters apply when the grant_type parameter for /as/token.oauth2 is set to urn:ietf:params:oauth:grant-type:token_exchange.

Parameter	Description
resource	An absolute URI that indicates the target service or resource where the client will use the requested security token. This information lets the authorization server apply the appropriate policy for the target.
	For example, the authorization server can ensure that the token it issues has the type, content, and encryption that the target requires. If the issued token will be used at multiple targets, multiple resource parameters can be used.
audience	The logical name of the target where the client will use the requested security token. This serves a purpose like the resource parameter. To interpret the audience, the value must be something that both the client and the authorization server understand. The value must be unique within the authorization server. If the issued token will be used at multiple targets, multiple audience parameters can be used. When indicating multiple targets, audience and resource parameters can be used together.
scope	The scope of the access request expressed as a list of space-delimited, case-sensitive strings. Scopes can also be constrained on a client-to-client basis. For more information, see <i>Scopes and scope management</i> on page 486.
requested_token_	typeAn identifier for the type of the requested security token. If the requested type is not specified, the authorization server can determine which type is required by the target that the resource or audience parameter identifies.

Parameter	Description
subject_token (Required)	A security token that represents the identity of the party on behalf of whom the request is being made. Typically, the subject of this token will be the subject of the security token issued in response to the request.
	Important:
	To avoid recording this parameter in web server logs, only pass it in the message body using the HTTP POST method.
<pre>subject_token_type (Required)</pre>	An identifier for the type of security token in the subject_token parameter.
actor_token	A security token that represents the identity of the acting party. Typically, this will be the party that is authorized to use the requested security token and act on behalf of the subject.
	Important:
	To avoid recording this parameter in web server logs, only pass it in the message body using the HTTP POST method.
actor_token_type	An identifier for the type of security token in the actor_token parameter.
	Important:
	This parameter is required when the actor_token parameter is present in the request but must not be included otherwise.
	_ · · · · · · · · · · · · · · · · · · ·

Access token validation grant type

The following parameter applies when the <code>grant_type</code> parameter for <code>/as/token.oauth2</code> is set to <code>urn:pingidentity.com:oauth2:grant type:validate bearer.</code>

Parameter	Description
token	The bearer access token to be validated.
(Required)	Important:
	To avoid recording this parameter in web server logs, only pass it in the message body using the HTTP POST method.

This validation grant type is a custom PingFederate OAuth extension that enables a resource server (RS) to communicate with the OAuth AS while leveraging the established communication and encoding patterns from OAuth 2.0. This grant type allows an RS to verify with the OAuth AS on the validity of a bearer access token that it has received from a client making a protected-resources call.



Client authentication is not required. In other words, when creating a client for the sole purpose of validating access tokens, the **Client Secret** field is optional. For this grant type, the RS acts in the role of a client for the request/response exchange with the OAuth AS to make the validation call.

The response is a standard OAuth access-token response from the token endpoint with some extensions and minor semantic differences in the treatment of some of the parameters. The returned token is in a JSON structure with name-to-value elements or name-to-array elements.

The token type is urn:pingidentity.com:oauth2:validated_token, a URN indicating the token represents the attributes associated with the validated access token passed on the request. A client_id element is returned indicating the client identifier of the client to whom the grant was made. A scope element is returned, if the scope is greater than the default implied scope, indicating the approved scope of the grant. If the issuing access token management (ATM) instance is configured to expand scope groups, the response includes the corresponding sub scopes instead of the scope groups. The expires_in element indicates for how many more seconds the token is valid; the value can increase on subsequent validation calls if a token lifetime extension policy is in place.

Sample response when scope group expansion is disabled (the default)

```
"access_token": {
    "Username": "joe",
    "OrgName": "Ping Identity Corporation"
},
"scope": "openid AAAGroup",
"token_type": "urn:pingidentity.com:oauth2:validated_token",
"expires_in": 7121,
"client_id": "ac_oic_client"
}
```

AAAGroup is a scope group.

Sample response when scope group expansion is enabled

```
"access_token": {
    "Username": "joe",
    "OrgName": "Ping Identity Corporation"
},
    "scope": "openid AAA1 AAA2",
    "token_type": "urn:pingidentity.com:oauth2:validated_token",
    "expires_in": 7169,
    "client_id": "ac_oic_client"
}
```

AAA1 and AAA2 are the expanded outcome of AAAGroup.

Validate against all eligible ATM instances

If multiple ATM instances are eligible, the configuration of the RS client determines whether it must specify the desired ATM instance in its token validation requests. For more information, see .

Once an ATM instance is chosen, PingFederate considers the per-instance session validation settings and processes the validation request. For more information, see .

Introspection endpoint

A resource server (RS) client uses the introspection endpoint to validate an access token or a refresh token prior to granting access to a protected-resources call.

The OAuth 2.0 Token Introspection documentation defines the introspection endpoint.



This endpoint accepts only the HTTP POST method.

Endpoint: /as/introspect.oauth2

When transmitting through the HTTP POST method, the required Content-Type value is application/x-www-form-urlencoded. The RS acts in the role of a client for the request/response exchange with the PingFederate OAuth AS. The validation call is made using the following parameters.

Parameter	Description
token	The bearer access token or refresh token to be validated.
(Required)	Important: To avoid recording this parameter in web server logs, only pass it in the message body using the HTTP POST method.
token_type_hint	A hint about the type of token submitted for validation. PingFederate supports values of access_token and refresh_token. Required only when validating a refresh token.

Client authentication is not required; when creating a client for the sole purpose of validating access tokens and refresh tokens, the Client Secret field is optional.

The response is in a JSON structure with a list of name-to-value elements. If a token is valid, the OAuth AS returns to the client a JSON object with the following elements:

- An "active": true element to indicate the token is valid. This is also the only element returned by the OAuth AS for a valid refresh token.
- A client id element to indicate the client identifier of the client to whom the grant was made.
- A scope element, if the scope is greater than the default implied scope, indicating the approved scope of the grant. If you configured the issuing access token management (ATM) instance to expand scope groups, the response includes the corresponding sub scopes instead of the scope groups.
- An exp element indicates the token is valid until the number of seconds since January 1, 1970 UTC (epoch time).
- Other elements from the access token.



Note:

The response includes the username and subject (sub) elements only if they were mapped in the ATM instance.

If a token is invalid, the OAuth AS returns { "active": false} to the client.

A sample response for a valid access token when scope group expansion is disabled (the default)

```
"scope": "openid AAAGroup",
```

```
"active": true,
"OrgName": "Ping Identity Corporation",
"token_type": "Bearer",
"exp": 1556823489,
"client id": "ac oic client"
```

AAAGroup is a scope group.

A sample response for a valid access token when scope group expansion is enabled

```
"scope": "openid AAA1 AAA2",
  "active": true,
  "OrgName": "Ping Identity Corporation",
  "token_type": "Bearer",
"exp": 1556823764,
  "client id": "ac oic client"
}
```

AAA1 and AAA2 are the expanded outcome of AAAGroup.

Response for a valid refresh token

```
"active": true
"exp": 1556823764
```



Note:

If the refresh token is configured to never expire, the "exp" attribute will not be returned.

Response for an invalid token

```
{"active":false}
```

OAuth client identification and authentication

The authentication requirement of this endpoint depends on the client authentication method configured for the clients.

Parameters
Clients can present their client identifier and client secret using the HTTP Basic authentication scheme, where the client identifier is the username, and the client secret is the password.
Clients can provide credentials using the request parameters client_id and client_secret.
Important:
This is a sensitive parameter. To avoid recording it in web server logs, only pass in this parameter with the HTTP POST method in the message body, or through the HTTP Basic authentication scheme.

Authentication method	Parameters
Client certificate	Clients must present their client certificate for mutual TLS authentication. The issuer and the subject distinguished name (DN) of the client certificate must match values configured for the clients.
Private key JWT	Clients must include request parameters client_assertion_type and client_assertion in the message body of their requests.
	client_assertion_type
	The value describes the format of the assertion as defined by the authorization server. For the private_key_jwt client authentication method, the value is urn:ietf:params:oauth:client-assertion-type:jwt-bearer.
	client_assertion
	The value is the authentication token.
	Example
	<pre>client_assertion_type= urn%3Aietf%3Aparams%3Aoauth% 3Aclient-assertion-type%3Ajwt-bearer& client_assertion= eyJhbGciOiJSUzI1NiIsLbSWi1YO-TILOd4L7ZCg&</pre>
	Note:
	For readability, line breaks are inserted and the authentication token is truncated.
	For more information about the private_key_jwt client authentication method, see Client Authentication and Using Assertions for Client Authentication.
None	Clients must pass in the client_id parameter in a query string or the message body to identify themselves.

OAuth access token management parameters

PingFederate supports multiple access token management (ATM) instances. Clients can specify an ATM instance by providing the ATM ID (access_token_manager_id) or a resource URI (aud) in their requests to the PingFederate OAuth AS.

Parameter	Description
access_token_manage	arThetaccess_token_manager_id value is the instance ID of the desired ATM instance. When specified, PingFederate uses the desired ATM instance for the request if it is eligible; otherwise it aborts the request.
	Note:
	When the access_token_manager_id parameter is specified, PingFederate ignores the aud parameter.

Parameter	Description
aud	The aud is the resource URI the client wants to access. The provided value is matched against resource URIs configured in access token management instances. When a match is found, PingFederate uses the corresponding access token management instance for the request if it is eligible; otherwise it aborts the request.

A match can be an exact match or a partial match where the provided URI has the same scheme and authority parts and a more specific path contained within the path of the pre-configured resource URI. PingFederate takes an exact match over a partial match. If there are multiple partial matches, PingFederate takes the partial match where the provided URI matches more specifically against the pre-configured resource URI.

Example

A partial match

A resource URI of https://app.example.local is a partial match for the following provided URIs:

- https://app.example.local/file1.ext
- https://app.example.local/path/file2.ext
- https://app.example.local/path/more

Example

An exact match is a better match than a partial match

Access Token Management instances	Resource URIs (configured)
ATM1	https://localhost:9031/app1
	https://localhost:9031/app2/data
	https://app.example.local
ATM2	https://localhost:9031/app1/data
	https://localhost:9031/app2/data/get

https://localhost:9031/app1 (a resource URI pre-configured for ATM1) is a partial match for https://localhost:9031/app1/data (the provided URI). However, PingFederate chooses ATM2 because https://localhost:9031/app1/data (a resource URI pre-configured for ATM2) is an exact match against the provided URI.

Example

A more specific partial match is a better match

Both https://localhost:9031/app2/data (a resource URI for ATM1) and https://localhost:9031/app2/data/get (a resource URI for ATM2) are partial matches for https://localhost:9031/app2/data/get/sample (the provided URI). However, PingFederate chooses ATM2 because https://localhost:9031/app2/data/get matches more specifically against the provided URI.

Validate against all eligible ATM instances

If multiple ATM instances are eligible, the configuration of the RS client determines whether it must specify the desired ATM instance in its token validation requests. For more information, see .

Once an ATM instance is chosen, PingFederate considers the per-instance session validation settings and processes the validation request. For more information, see .

Token revocation endpoint

The token revocation endpoint allows clients to notify the authorization server that a previously obtained refresh or access token is no longer needed. The revocation request invalidates the actual token and possibly other tokens based on the same authorization grant.

The OAuth 2.0 Token Revocation documentation defines the token revocation endpoint.



This endpoint accepts only the HTTP POST method.

Endpoint: /as/revoke token.oauth2



Important:

Only Internally Managed Reference Tokens support direct access token revocation. JSON web token (JWT) type access tokens do not support direct revocation. JWT access tokens can only be indirectly revoked if the associated refresh token is revoked, and the JWT's configuration field Access Grant GUID **Claim Name** is set for the given access token manager instance.

When the authorization server revokes a refresh token, it also revokes the associated access grant and access tokens. When the authorization server revokes an access token, the associated access grant and refresh token remain untouched with the exception of the implicit grant type. If the Reuse Existing Persistent Access Grants for GrantTypes check box is selected in the System # OAuth Settings # **Authorization Server Settings** window, the implicit access grant will also be revoked with the access token.

The following table describes parameters for this endpoint. The required Content-Type value is application/x-www-form-urlencoded when transmitting through the HTTP POST method.

Parameter	Description
token	The token that the client wants to revoke.
(Required)	Important:
	To avoid recording this parameter in web server logs, only pass it in the message body using the HTTP POST method.
token_type_hint	A hint about the type of token submitted for revocation. PingFederate supports values of access_token and refresh_token.

The following table describes parameters for this endpoint. The required Content-Type value is application/x-www-form-urlencoded.

OAuth client identification and authentication

The authentication requirement of this endpoint depends on the client authentication method configured for the clients.

Authentication method	Parameters
Client secret	Clients can present their client identifier and client secret using the HTTP Basic authentication scheme, where the client identifier is the username, and the client secret is the password.
	Clients can provide credentials using the request parameters client_id and client_secret.
	Important:
	This is a sensitive parameter. To avoid recording it in web server logs, only pass in this parameter with the HTTP POST method in the message body, or through the HTTP Basic authentication scheme.
Client certificate	Clients must present their client certificate for mutual TLS authentication. The issuer and the subject distinguished name (DN) of the client certificate must match values configured for the clients.
Private key JWT	Clients must include request parameters client_assertion_type and client_assertion in the message body of their requests.
	client_assertion_type
	The value describes the format of the assertion as defined by the authorization server. For the private_key_jwt client authentication method, the value is urn:ietf:params:oauth:client-assertion-type:jwt-bearer.
	client_assertion
	The value is the authentication token.
	Example
	client_assertion_type= urn%3Aietf%3Aparams%3Aoauth% 3Aclient-assertion-type%3Ajwt-bearer& client_assertion= eyJhbGciOiJSUzI1NiIsLbSWi1YO-TILOd4L7ZCg&
	Note: For readability, line breaks are inserted and the authentication token is truncated.
	For more information about the private_key_jwt client authentication method, see Client Authentication and Using Assertions for Client Authentication.
None	Clients must pass in the client_id parameter in a query string or the message body to identify themselves.

Grant-management endpoint

Resource owners use the grant-management endpoint to view, and optionally revoke, the persistent access grants they have made.

Two grant-management endpoints are provided. One is for use with parameters. This endpoint is not part of the OAuth specification, but many OAuth providers offer a similar function.

Grants associated with the USER KEY of the authenticated user are displayed. The same attribute mappings from the authentication source to USER KEY, which are used for the authorization endpoint, are used here to look up the user's existing grants.

Endpoints: /as/grants.oauth2 and /as/oauth access grants.ping

The following table describes the available parameters for the /as/grants.oauth2 endpoint. Use only one of them as needed.

Parameter	Description
idp Or PartnerIdpId	Indicates the entity ID of the connection ID of the identity provider (IdP) with whom to initiate browser single sign-on (SSO) for user authentication.
pfidpadapterid	Indicates the IdP adapter instance ID of the adapter to use for user authentication.
	Note: This parameter may be overridden by policy based on authentication selection configuration. For example, the OAuth Scope Authentication Selector could enforce the use of a given adapter based on client-requested scopes.

If no recent user attributes are found for the session context, the user is redirected to /as/ oauth access grants.ping to initiate the authentication process, which behaves in the same way as the authorization endpoint.

Dynamic client registration endpoint

The client registration endpoint allows developers to dynamically register OAuth clients on a PingFederate authorization server.

The OAuth 2.0 Dynamic Client Registration Protocol defines this endpoint. Developers can send client registrations with the desired properties, such as client metadata, to this endpoint. If the requests are valid, PingFederate evaluates them and returns a response with a client ID and the registered client metadata values.



Note:

This runtime endpoint is only active when the dynamic registration client is enabled and configured.



Important:

As dynamic client registration can expose your server to unwanted client registrations, we recommend protecting PingFederate by requiring an initial access token, configuring one or more client registration policies, and protecting access to the dynamic client registration endpoint.

You can configure access token requirement and client registration policies using the System # OAuth Settings # Client Settings window. To further protect against unauthorized access to the dynamic client registration endpoint, consider using PingAccess or your choice of web access management solution to do SO.



This endpoint accepts only the HTTP POST method.

Endpoint: /as/clients.oauth2

Both the request and the response follow the OAuth 2.0 Dynamic Client Registration Protocol.

Example

Example 1

A developer wants to register a client that supports the authorization code flow, two redirection URIs, two scopes, and HTTP Basic as the client authentication method. In this example, PingFederate is not configured to require an initial access token.

Request

```
POST /as/clients.oauth2 HTTP/1.1
Content-Type: application/json
Accept: application/json
Host: sso.example.com
  "client name": "Example Org Sample One",
  "redirect uris":[
    "https://example.org/app1",
    "https://example.org/appM"
  "scope": "email phone",
  "grant types":[
    "authorization code"
  ]
}
```

Response

```
HTTP/1.1 201 Created
Date: Fri, 13 Oct 2017 12:34:56 GMT
Referrer-Policy: origin
Content-Type: application/json
Transfer-Encoding: chunked
  "client id": "dc-F3JxcBlNCtjk36J3Yi4yQK",
  "client name": "Example Org Sample One",
  "redirect uris": [
    "https://example.org/app1",
    "https://example.org/appM"
  "token endpoint auth_method": "client_secret_basic",
  "grant types": [
    "authorization code"
```

```
"client_secret": "fYhGUjnkjGp0UPQGaAfdcS",
   "client_secret_expires_at": 0,
   "scope": "phone email",
   "validate_using_all_eligible_atms": false,
   "refresh_token_rolling_policy": "server_default",
   "persistent_grant_expiration_type": "server_default",
   "grant_access_session_revocation_api": false
   "grant_access_session_management_api": false
}
```

PingFederate returns 201 Created, the client ID, and other registered client metadata after creating the new client.

Additionally, when a registration request does not specify a client authentication method (token_endpoint_auth_method), PingFederate defaults to client_secret_basic per OAuth 2.0 Dynamic Client Registration Protocol.

Example

Example 2

A developer wants to register a client that supports the authorization code flow, refresh tokens, one redirection URI, one profile scope, and HTTP Basic as the client authentication method. In this example, PingFederate is not configured to require an initial access token. However, the profile scope is restricted. As a result, the registration request should fail.

Request

```
POST /as/clients.oauth2 HTTP/1.1
Content-Type: application/json
Accept: application/json
Host: www.example.com

{
    "client_name":"Example Org Sample Two",
    "redirect_uris":[
        "https://example.org/app2"
    ],
    "scope":"profile",
    "grant_types":[
        "authorization_code",
        "refresh_token"
    ]
}
```

Response

```
HTTP/1.1 400 Bad Request
Date: Fri, 13 Oct 2017 13:00:00 GMT
Referrer-Policy: origin
Content-Type: application/json
Transfer-Encoding: chunked
{
    "error": "invalid_client_metadata",
    "error_description": "The requested scope is invalid."
}
```

PingFederate returns 400 Bad Request and the relevant error message when a client registration fails.

Example

Example 3

A developer wants to register a client that supports the authorization code flow, two redirection URIs, two scopes, and HTTP Basic as the client authentication method. In this example, PingFederate is configured to require an initial access token.

Request

```
POST /as/clients.oauth2 HTTP/1.1
Content-Type: application/json
Accept: application/json
Authorization: Bearer
eyJhbGciOiJSUzI1NiIsImtpZCI6ImsxIn0.eyJzY29wZSI6WyJkQ1IiXSwiY2xpZW50X21kX25hbWUiOi
CHtcQ79Wefz2Sw5GOB5LfV9mWJ0n3vzJ93Ie7wbEAkalIFg53J-9e7s59MjA1igx6ybf1GMQ9QAjYobs-
jM24arJZZgopEXvcx6IQpyU8U4AMTJ7tr9Lmody8P0QZOKcUDBTT5egv9vr5NuXCtUBfVPhGZ-3p5g5mwrn
oDhmilbmiga4319YSFfX5-
U31i9XPeN3JZB2ukLbTFjjVIVLJIInbSR IFTWP5Irg92aXLrIfm5MvBp8D1f0U6xYjbgjvw9QKNiFFVD7o
Host: www.example.com
  "client name": "Example Org Sample Three",
  "redirect uris":[
    "https://example.org/app3",
    "https://example.org/appN"
  "scope": "email phone",
  "grant types":[
    "authorization code"
}
```

Response

```
HTTP/1.1 201 Created
Date: Fri, 13 Oct 2017 15:30:00 GMT
Referrer-Policy: origin
Content-Type: application/json
Transfer-Encoding: chunked
  "client id": "dc-rqUtii4vRXj5NMztkAeJ1S",
  "client name": "Example Org Sample Three",
  "redirect uris": [
    "https://example.org/app3",
    "https://example.org/appN"
  "token endpoint auth method": "client_secret_basic",
  "grant types": [
    "authorization code"
  "client secret": "p7MD0Ul1DNI9xRDc5kc0xs",
  "client secret expires at": 0,
  "scope": "phone email"
  "validate using all eligible_atms": false,
  "refresh token rolling policy": "server default",
  "persistent grant expiration type": "server default",
  "grant access session revocation api": false
  "grant access session management api": false
```

The registration request must include an Authorization HTTP header with a valid access token as its

If the authorization fails, PingFederate returns the following JSON payload in the response.

```
"error": "invalid access token",
 "error description": "Please provide a valid Access Token with the correct
scope"
```

Device authorization endpoint

The device authorization endpoint allows a user to grant authorization to a device client using a browser on a second device, such as a smart phone or a computer.

The OAuth 2.0 Device Authorization Grant defines the device authorization endpoint. Based on the specification, the device sends a device authorization request to PingFederate, the authorization server (AS), at its device authorization endpoint.



Note:

Per OAuth specifications, this endpoint accepts only the HTTP POST method.

Endpoint: /as/device_authz.oauth2

The required Content-Type value is application/x-www-form-urlencoded. The following table describes parameters for this endpoint.

Parameter	Description
client_id	A unique identifier the client provides to the resource server (RS) to identify itself. This identifier is included with every request the client makes
scope	The scope of the access request expressed as a list of space-delimited, case-
(Optional)	sensitive strings.
	Scopes can also be constrained on a client-to-client basis. For more information about scopes, see <i>Scopes and scope management</i> on page 486.

Both the request and the response follow the OAuth 2.0 Device Authorization Grant.

Example

Example request

```
POST /as/device authz.oauth2 HTTP/1.1
Content-Type: application/x-www-form-urlencoded
Host: www.example.com
client id=df client
```

Example

Response codes and example responses

200 - Success

```
HTTP/1.1 200 OK

...
{
    "user_code": "YYD6-CD4T",
    "device_code": "4EHsIngavzIPvvqMlFgQlseTCsH7EpU75f9yGvj60T",
    "interval": 5,
    "verification_uri_complete": "https://www.example.com/as/
user_authz.oauth2?user_code=YYD6-CD4T",
    "verification_uri": "https://www.example.com/as/user_authz.oauth2",
    "expires_in": 600
}
```

400 - Bad Request

```
HTTP/1.1 400 Bad Request
...
{
    "error_description": "The requested scope(s) must be blank or a subset
    of the provided scopes.",
        "error": "invalid_scope"
}
```

401 - Unauthorized

```
HTTP/1.1 401 Unauthorized
...
{
    "error_description": "Invalid client or client credentials.",
    "error": "invalid_client"
}
```

OAuth client identification and authentication

The authentication requirement of this endpoint depends on the client authentication method configured for the clients.

Authentication method	Parameters
Client secret	Clients can present their client identifier and client secret using the HTTP Basic authentication scheme, where the client identifier is the username, and the client secret is the password.
	Clients can provide credentials using the request parameters client_id and client_secret.
	Important:
	This is a sensitive parameter. To avoid recording it in web server logs, only pass in this parameter with the HTTP POST method in the message body, or through the HTTP Basic authentication scheme.
Client certificate	Clients must present their client certificate for mutual TLS authentication. The issuer and the subject distinguished name (DN) of the client certificate must match values configured for the clients.
Private key JWT	Clients must include request parameters client_assertion_type and client_assertion in the message body of their requests.
	client_assertion_type
	The value describes the format of the assertion as defined by the authorization server. For the private_key_jwt client authentication method, the value is urn:ietf:params:oauth:client-assertion-type:jwt-bearer.
	client_assertion
	The value is the authentication token.
	Example
	<pre>client_assertion_type= urn%3Aietf%3Aparams%3Aoauth% 3Aclient-assertion-type%3Ajwt-bearer& client_assertion= eyJhbGciOiJSUzI1NiIsLbSWi1YO-TILOd4L7ZCg&</pre>
	Note: For readability, line breaks are inserted and the authentication token is truncated.
	For more information about the private_key_jwt client authentication method, see Client Authentication and Using Assertions for Client Authentication.
None	Clients must pass in the client_id parameter in a query string or the message body to identify themselves.

User authorization endpoint

The user authorization endpoint allows a user to grant authorization to a device client using a browser on a second device, such as a smart phone or a computer.

Based on the OAuth 2.0 Device Authorization Grant specification, the user goes to the user authorization endpoint of the PingFederate authorization server (AS) to complete the authorization process.



Note:

This endpoint accepts the HTTP GET and POST methods.

Endpoint: /as/user authz.oauth2

The following table describes parameter for this endpoint. The required Content-Type value is application/x-www-form-urlencoded when transmitting through the HTTP POST method.

Parameter	Description
user_code	This value represents the activation code.
(Optional)	

Both the request and the response follow the OAuth 2.0 Device Authorization Grant.

Example

Example request

```
POST /as/user authz.oauth2 HTTP/1.1
Content-Type: application/x-www-form-urlencoded
Host: www.example.com
user code=QQWP-TJ6B
```

Subsequent responses

Phase 1: Activation code verification

If the verification request does not include an activation code, PingFederate returns the Connect a device (user code prompt) page, prompting the user to enter the activation code shown by the device.

If the verification request includes an activation code, PingFederate returns the Connect a device (pre-populated user code prompt) page, prompting the user to confirm the activation code from the verification request matches the activation code shown by the device. PingFederate skips this step if the Bypass Activation Code Confirmation option is enabled globally or individually for that invoking client.

PingFederate validates the activation code, prompts the user to enter another activation code if it is invalid, or moves to the next phase.

Phase 2: Authentication

PingFederate prompts the user to fulfill the authentication requirements based on OAuth grant mapping configurations and authentication policies.

If the user fulfills the authentication requirements, PingFederate moves to the next phase; otherwise it returns an error message to the user.

PingFederate returns the **Request for Approval** page, prompting the user to approve or deny the requested scopes. PingFederate skips this step if the **Bypass Authorization Approval** option is enabled globally or individually for that invoking client and the user has granted authorization for the requested scopes previously.

PingFederate returns the **Connect a device (result)** page to the user. The message reflects the authorization status.

If the user approves the requested scopes, the next time the device sends a device access token request to PingFederate at its token endpoint, PingFederate returns an access token to the device.

When an error occurs, PingFederate returns 400 Bad Request in response to the device access token request.

Example

Examples of 400 Bad Request

```
HTTP/1.1 400 Bad Request

...
{"error_description":"Authorization request is denied","error":"access_denied"}

HTTP/1.1 400 Bad Request

...
{"error_description":"Device code not found, expired or invalid","error":"invalid_grant"}

HTTP/1.1 400 Bad Request

...
{"error_description":"The authorization request has expired.","error":"expired_token"}
```

OpenID Provider configuration endpoint

The OpenID Provider (OP) configuration endpoint provides configuration information for the OAuth clients to interface with PingFederate using the OpenID Connect protocol.

This endpoint returns configuration information that is controlled by a template file and can be customized to suit multiple use cases simultaneously.

This public endpoint accepts HTTP GET requests without authentication.

Endpoint: /.well-known/openid-configuration

The following table describes the parameter for this endpoint.

Parameter	Description
policy_id	Indicates the OpenID Connect policy from which PingFederate derives the attributes to include under claims_supported in the response body.
	If omitted, PingFederate includes the attributes based on the default policy.

Example

Example response

```
$ curl -s https://localhost:9031/.well-known/openid-configuration|python -m
json.tool
    "authorization endpoint": "https://localhost:9031/as/
authorization.oauth2",
    "backchannel authentication endpoint": "https://localhost:9031/as/bc-
auth.ciba",
    "backchannel authentication request signing alg values supported": [
        "RS256",
        "RS384",
        "RS512",
        "ES256",
        "ES384",
        "ES512",
        "PS256",
        "PS384",
        "PS512"
    "backchannel token_delivery_modes_supported": [
        "poll",
        "ping"
    "backchannel_user_code_parameter_supported": true,
    "code challenge methods supported": [
        "plain",
        "S256"
    "claim_types_supported": [
        "normal"
    "claims parameter_supported": false,
    "claims supported": [
        "address",
        "birthdate",
        "email",
        "email_verified",
        "family_name",
        "gender",
        "given name",
        "locale",
        "middle name",
        "name",
        "nickname",
        "phone number",
        "phone number_verified",
        "picture",
        "preferred_username",
        "profile",
        "sub",
        "updated_at",
        "website",
        "zoneinfo"
    "device authorization endpoint": "https://localhost:9031/as/
device authz.oauth2",
    "grant types_supported": [
        "implicit",
        "authorization code",
```

```
"refresh token",
        "password",
        "client credentials",
        "urn:pingidentity.com:oauth2:grant type:validate bearer",
        "urn:ietf:params:oauth:grant-type:jwt-bearer",
        "urn:ietf:params:oauth:grant-type:saml2-bearer",
        "urn:ietf:params:oauth:grant-type:device_code",
        "urn:openid:params:grant-type:ciba"
    "id token_encryption_alg_values_supported": [
        "dir",
        "A128KW"
        "A192KW",
        "A256KW",
        "A128GCMKW",
        "A192GCMKW"
        "A256GCMKW"
        "ECDH-ES",
        "ECDH-ES+A128KW",
        "ECDH-ES+A192KW"
        "ECDH-ES+A256KW",
        "RSA-OAEP"
    "id_token_encryption_enc_values_supported": [
        "A128CBC-HS256",
        "A192CBC-HS384"
        "A256CBC-HS512",
        "A128GCM",
        "A192GCM"
        "A256GCM"
    "id token signing_alg_values_supported": [
        "none"
        "HS256",
        "HS384",
        "HS512",
        "RS256",
        "RS384",
        "RS512",
        "ES256",
        "ES384",
        "ES512",
        "PS256",
        "PS384",
        "PS512"
    "introspection endpoint": "https://localhost:9031/as/introspect.oauth2",
    "issuer": "https://localhost:9031",
    "jwks_uri": "https://localhost:9031/pf/JWKS",
    "ping_end_session_endpoint": "https://localhost:9031/idp/startSLO.ping", "ping_revoked_sris_endpoint": "https://localhost:9031/pf-ws/rest/
sessionMgmt/revokedSris",
    "registration endpoint": "https://localhost:9031/as/clients.oauth2",
    "request object signing alg values supported": [
        "RS256",
        "RS384"
        "RS512",
        "ES256"
        "ES384"
        "ES512"
        "PS256",
        "PS384"
        "PS512"
    ],
```

```
"request parameter supported": true,
"request uri parameter supported": false,
"response_modes_supported": [
    "fragment",
    "query",
    "form_post"
"response types supported": [
    "code",
    "token",
    "id token",
    "code token",
    "code id token",
    "token i\bar{d} token",
    "code token id_token"
"revocation endpoint": "https://localhost:9031/as/revoke_token.oauth2",
"scopes_supported": [
    "address",
    "phone",
    "edit",
    "openid"
    "profile",
    "admin",
    "email"
"subject types_supported": [
    "public",
    "pairwise"
"token endpoint": "https://www.example.com:9031/as/token.oauth2",
"token endpoint auth_methods_supported": [
    "client_secret_basic",
"client_secret_post",
    "private key jwt"
"token endpoint auth signing alg values_supported": [
    "RS256",
    "RS384",
    "RS512",
    "ES256",
    "ES384",
    "ES512",
    "PS256",
    "PS384",
    "PS512"
"userinfo endpoint": "https://localhost:9031/idp/userinfo.openid"
```

Notable metadata parameters

CIBA user code support

The backchannel_user_code_parameter_supported parameter indicates whether the default CIBA request policy supports user codes, which are an optional feature in the CIBA specification.

In the previous example, because the **User Code PCV** field is configured with a Password Credential Validator instance in the default CIBA request policy, the value of the **backchannel_user_code_parameter_supported** parameter is true. For more information,

see OpenID Connect Client Initiated Backchannel Authentication Flow and Defining a request policy on page 582.

Digital signature algorithms

```
The backchannel_authentication_request_signing_alg_values_supported, id_token_signing_alg_values_supported, token_endpoint_auth_signing_alg_values_supported, and request_object_signing_alg_values_supported parameters provide lists of supported algorithms to process digital signatures.
```

In this example, because PingFederate is integrated with a hardware security module (HSM) and configured to use static keys for OAuth and OpenID Connect, the endpoint includes additional RSASSA-PSS digital signature algorithms (PS256, PS384, and PS512) in its response. For more information on HSM integration and static keys, see *Supported hardware security modules* on page 169 and *Keys for OAuth and OpenID Connect* on page 616, respectively. Deploying PingFederate to run on a Java 11 environment will have the same result.

JWKS endpoint

The JWKS endpoint, jwks_uri, returns a set of public keys for OAuth and OpenID Connect. Clients can use this information to verify the integrity of asymmetrically-signed ID tokens, JSON web tokens (JWTs) for client authentication, and OpenID Connect request objects.

Scopes

The OP configuration endpoint returns all common static scopes and common scope groups but not exclusive static scopes, exclusive scope groups, common dynamic scopes, or exclusive dynamic scopes by default. The response can be customized by editing a template file to include or exclude individual scopes and scope groups.

Token endpoint

The token endpoint, token_endpoint, is used by clients to obtain access tokens and refresh tokens if applicable.

In the previous example, because the **Token Endpoint Base URL** is set to https://www.example.com:9031 in the **System # OAuth Settings # Authorization Server Settings** window, the **token_endpoint** value is set to https://www.example.com:9031/as/token.oauth2. For more information, see *Configuring authorization server settings* on page 475 and *Token endpoint* on page 1057.

UserInfo endpoint

OAuth clients can present access tokens to the UserInfo endpoint to retrieve additional information about the resource owners.

You can customize the amount of information presented by the endpoint by using OpenID Connect policies. Information can include specification-defined attributes (standard attributes) and non-standard attributes. Scopes, authorized by the users, also determine the attributes to be returned.

This endpoint accepts HTTP GET requests without parameters. Clients must present valid access tokens for authentication.

Example

Endpoint example: /idp/userinfo.openid

```
$ curl -s https://localhost:9031/idp/userinfo.openid -H 'Authorization:
Bearer eyJ...9-g'|python -m json.tool
{
    "email": "auser@example.com",
    "phone_number": "(555) 555-5555",
```



The self-contained access token in the Authorization HTTP header is truncated for readability.

If the access token presented is not valid, PingFederate returns 401 Unauthorized.

Self-contained tokens

If clients using self-contained access tokens are expected to contact the UserInfo endpoint, consider the following implications.

Client ID Claim Name

This field's default value is client id. When this field is configured with a value, PingFederate includes the client ID of the requesting client as a claim in the self-contained tokens. The claim name is the value of the Client ID Claim Name field.

If the field value is empty, PingFederate will not include the client ID of the requesting client in the self-contained tokens. In this scenario, the access token manager (ATM) instance used by the default OpenID Connect policy must remain accessible to all clients, or clients using self-contained access tokens issued by this ATM instance will not be able to retrieve additional claims from the UserInfo endpoint. Instead, they receive an HTTP status code 401 Unauthorized from PingFederate. For more information, see *Defining access control* on page 561.

Scope Claim Name

This field's default value is scope. When this field is configured with a value, PingFederate includes the requested scopes as a claim in the self-contained tokens. The claim name is the value of the Scope Claim Name field.

If the field value is empty, PingFederate will not include any scope information in the self-contained token, and clients using self-contained access tokens issued by this ATM instance will not be able to retrieve additional claims from the UserInfo endpoint. Instead, they receive an HTTP status code 403 Forbidden from PingFederate.

Web service interfaces and APIs

PingFederate provides two built-in, SOAP-accessible web services related to browser-based single sign-on (SSO).

These services can be used by client applications to manage partner connections and support integration of web applications, respectively.

Connection Management Service

The Connection Management Service enables creation and deletion of single connection configurations in PingFederate. This service can be used to migrate connections from one server environment to another, for example, from testing or staging to production, or to create new connections in a single server programmatically.



Z Tip:

PingFederate provides a command-line utility that can be used to export and modify connections, as well as other administrative-console configurations, and then import them to target environments. For more information, see *Automating configuration migration* on page 790.

SSO Directory Service

The SSO Directory Service provides web application developers with information regarding partner connections and adapter instances.



Applications accessing the Connection Management Service must first authenticate themselves to the PingFederate server. SSO Directory Service authentication is optional by default, but might be required. For more information, see Configuring service authentication on page 636.

Additionally, PingFederate provides REST-based web services and APIs for a variety of administrative and runtime tasks.

OAuth Client Management Service

A runtime API to manage OAuth client applications.

OAuth Access Grant Management Service

A runtime API to retrieve and revoke persistent grants. This API is intended for administrators to manage grants per client or per user.

OAuth Persistence Grant Management API

Another runtime API to retrieve and revoke persistent grants. This API is intended for the use case where clients can assume the responsibility of grant management, provided that the users authorize the clients to do so.

Session Revocation API

A runtime API allowing clients supporting the OpenID Connect protocol to guery revocation status of their sessions and add user sessions to the revocation list.

Administrative API

An administrative API to manage various PingFederate settings.

Connection Management Service

The Connection Management Service supports basic connection management capabilities and is accessible only on a PingFederate server running the administrative console.

The Connection Management Service is useful in a variety of circumstances. Consider the following use cases:

- Using the Connection Management Service as a utility, you can migrate changes to a partner connection through staging environments. For example, development, test, and production.
 - Using the Connection Management Service, you might need to make changes to URLs and keys to make the connection appropriate to the next environment.
- Using the Connection Management Service, an external application can update or delete connections programmatically, or create new ones using an exported connection XML file as a template.

You can find the WAR file for this service, pf-mgmt-ws.war, in the <pf install>/pingfederate/ server/default/deploy2 directory.

Note:

If you do not want to allow use of the service, do not deploy it: remove the WAR file from the deploy2 directory.

The SOAP-accessible service endpoint is pf-mgmt-ws/ws/ConnectionMigrationMgr.

The web services Description Language (WSDL) document describing this service can be retrieved from / pf-mgmt-ws/ws/ConnectionMigrationMgr?wsdl.

Exporting a connection

You can export a connection either manually, using the administrative console, or programmatically through a call to the Connection Management Service.

Whether you export a connection manually or programmatically, the exported XML complies with the standard SAML 2.0 metadata format, with extensions to capture PingFederate's proprietary configuration. Most connection configuration information is contained in the XML markup, with the exception of global configuration items such as adapter instances, datastores, and key pairs. Adapter instances and datastores are referenced by ID, and key pairs are referenced by the MD5 fingerprint of their X.509 certificate. Public certificates, such as the partner's signature verification certificate, are included completely (base-64 encoded).

Export manually

For information about using the administrative console to export connections, see Accessing SP connections on page 399 or Accessing IdP connections on page 662.

Export via the Connection Management Service

The Connection Management Service exposes the following method for exporting connections.

```
public string getConnection( String entityId, String role,) throws
 IOException
```

The entityId parameter is the connection ID, which identifies the connection to be deleted. The role parameter is the connection role, the identity provider (IdP) or the service provider (SP).

Code sample

The following example invokes this web service to export a connection.

```
Service service = new Service();
Call call = (Call) service.createCall();
call.setUsername("username");
call.setPassword("password");
call.setTargetEndpointAddress("https://localhost:9999/pf-mgmt-ws/ws/
ConnectionMigrationMgr");
call.setOperationName("getConnection");
Object result = call.invoke(new Object[] {"<entityId>", "SP"});
```

Importing connections

Moving a connection from one PingFederate server to another requires care, as the target server must contain the global configuration items, such as datastores, key pairs, and adapter instances, that the connection references.

Changing the references in the XML file, either manually or programmatically, may be necessary to adjust the connection to the target PingFederate environment.

Once the required changes are made to the XML file, developers can use the Connection Management Service to import the connection into a different instance of PingFederate.



Alternatively, you can import XML connection files through the PingFederate administrative console. For more information, see Accessing SP connections on page 399 or Accessing IdP connections on page 662. You can also import the connections into PingFederate manually by copying them into the <pf install>/pingfederate/server/default/ data/connection-deployer directory.

PingFederate scans this directory periodically and imports connections automatically.



CAUTION:

Manually importing a connection always overwrites an existing connection with the same ID. The web service provides a switch to disallow this behavior, if desired. For more information, see below.

The web service exposes the following method for importing connections.

```
public void saveConnection( String xml, boolean allowUpdate) throws
 IOException
```

The xml parameter is the complete representation of the connection retrieved by your application from an exported connection file, and optionally modified.

If allowUpdate is false, the web service can be used only to add a new connection. An error occurs if a connection already exists with the same connection ID and federation protocol in the XML. If allowUpdate is true and the connection already exists, it will be overwritten.

Sample code

The following example uses the Apache AXIS libraries to invoke this web service to create a new connection.

```
Service service = new Service();
    Call call = (Call) service.createCall();
    call.setUsername("username");
    call.setPassword("password");
    String addr = "https://localhost:9999/pf-mgmt-ws/ws/
ConnectionMigrationMgr";
    call.setTargetEndpointAddress(addr);
    call.setOperationName("saveConnection");
   String xml = "<EntityDescriptor entityID=\"some entity id\"</pre>
  </EntityDescriptor>";
   boolean allowUpdate = false;
    call.invoke(new Object[]{xml, allowUpdate});
```

Deleting connections

You can invoke the web service to delete connections.

The web service exposes the following method for connection deletion.

```
public void deleteConnection( String entityId, String role) throws
 IOException
```

The entityId parameter is the connection ID, which identifies the connection to be deleted. The role parameter is the connection role, identity provider (IdP) or service provider (SP).

Code sample

The following example uses the Apache AXIS libraries to invoke this web service to delete a connection.

```
Service service = new Service();
Call call = (Call) service.createCall();
call.setUsername("username");
call.setPassword("password");
call.setTargetEndpointAddress(
   "https://localhost:9999/pf-mgmt-ws/ws/ConnectionMigrationMgr"
   );
call.setOperationName("deleteConnection");
call.invoke(new Object[]{"entityid", "SP"});
```

Cluster configuration replication

A web service endpoint is available to replicate the administrative-console configuration to other nodes in a PingFederate cluster from the Connection Management Service.

The cluster configuration replication web service endpoint allows a client of this web service to create, update, or delete a connection, and then push the new configuration to the other cluster nodes.

The service endpoint is /pf-mgmt-ws/ws/ConfigReplication.

The WSDL document describing this service can be retrieved from /pf-mgmt-ws/ws/ConfigReplication?wsdl.

The web service exposes the following method: public void replicateConfiguration();.

Code sample

Below is example client code using the Apache AXIS libraries that invokes the configuration replication functionality.

```
Call call2 = (Call) service.createCall();
   call2.setUsername("joe");
   call2.setPassword("test");
   String addr2 = "https://localhost:9999/pf-mgmt-ws/ws/ConfigReplication";
   call2.setTargetEndpointAddress(addr2);
   call2.setOperationName("replicateConfiguration");
   call2.invoke(new Object[]{});
```

Validation disclaimer

Use the administrative console whenever possible to reduce the risk to data integrity.

The import process is not subject to the same rigorous data validation performed by the administrative user interface. Although some checks are made, it is possible to create invalid connections using the connection-migration process. As the XML is complex and validation is limited, use the administrative console to create the initial XML connection instead of attempting to create an XML connection from scratch. That way, changes necessary to the exported connection's XML representation can be held to a minimum, reducing the risk of compromising data integrity.

SSO Directory Service

PingFederate single sign-on (SSO) Directory Service allows applications to retrieve configuration data from a runtime PingFederate server. A PingFederate server in a cluster configured as an administrative console does not support this web service.

The SSO Directory Service lets web applications avoid storing and maintaining the data locally. These applications can retrieve the following types of data:

- A list of identity provider (IdP) partners
- A list of service provider (SP) partners

A list of SP adapter instances

The SSO Directory Service provides useful information for integrating an application with a PingFederate server. It is a way for the application to dynamically determine which partners can be used for SSO. This means applications do not need to be modified when new partners are configured in PingFederate.

You can find the WAR file for this module, pf-ws.war, in the pingfederate/server/default/ deploy directory.



If you do not want to allow use of the service, remove the WAR file from the deploy directory.

The service endpoint is: pf-ws/services/SSODirectoryService.

You can retrieve the WSDL document describing this service from /pf-ws/services/ SSODIrectoryService?wsdl.

You can retrieve a list using any of the following methods:

- getIDPList returns a list of active IdP connections configured for SP-initiated SSO. The list contains each IdP's connection ID and connection name.
- getSPList returns a list of active SP connections configured for IdP-initiated SSO. The list contains each SP's connection ID and connection name.



Note:

For either IdP or SP lists, connection IDs are returned as values for the XML tag <entityId>. Connection Names are returned as values for the XML tag <company>. For more information see SOAP request and response examples on page 1092.

- getAdapterInstanceList returns a list of SP adapter instances containing an ID and name.
- getIdpAdapterInstanceList returns a list of IdP adapter instances containing an ID and name.



Note:

These methods do not require input parameters.

The service is also available over HTTP. The query string for retrieving any of the lists is /pf-ws/ services/SSODirectoryService?method=<method name>.

Coding example

When you integrate a web application with PingFederate, use the single sign-on (SSO) Directory Service to generate a connection or adapter list. The code needed to create any of the lists is similar.

The following Java code example retrieves an identity provider (IdP) list from the web service. The program calls the getIDPList method in the SSO Directory Service to retrieve an IdP list and print it to the console. This example uses the Apache Axis library and includes optional code for authentication to the PingFederate server. For more information see Configuring service authentication on page 636. HTTPS is recommended when including credentials.

```
import org.apache.axis.client.Call;
import org.apache.axis.client.Service;
import java.net.URL;
import javax.xml.namespace.QName;
import com.pingidentity.ws.SSOEntity;
public class SSODirectoryClientSample
```

```
public static void main(String[] args) throws Exception
   Service service = new Service();
   Call call = (Call) service.createCall();
   call.setUsername("username");
    call.setPassword("pass");
   URL serviceUrl = new URL(
       "https://localhost:9031/pf-ws/services/
        SSODirectoryService");
   QName qn = new QName("urn:BeanService", "SSOEntity");
   call.registerTypeMapping(SSOEntity.class, qn,
      new org.apache.axis.encoding.ser.BeanSerializerFactory(
          SSOEntity.class, qn),
       new org.apache.axis.encoding.ser.BeanDeserializerFactory(
           SSOEntity.class, qn));
    call.setTargetEndpointAddress( serviceUrl );
   call.setOperationName( new QName(
      "http://www.pingidentity.com/servicesSSODirectoryService",
        "getIDPList"));
   Object result = call.invoke( new Object[] {} );
    if (result instanceof SSOEntity[])
       SSOEntity[] idpArray = (SSOEntity[]) result;
       for (SSOEntity idp : idpArray)
        System.out.println(idp.getEntityId() + " " +
          idp.getCompany());
   else
       System.out.println("Received problem response from
          server: " + result);
 }
}
```

SOAP request and response examples

A client application must send a SOAP request to the PingFederate server specifying the requested web service and the specific method.

The following is a typical SOAP request for an identity provider (IdP) list using the single sign-on (SSO) Directory Service.

The PingFederate server's web service returns a response containing the list you requested. The following is an example of a typical SOAP response for an IdP list.

```
<?xml version="1.0" encoding="UTF-8" ?>
<soapenv:Envelope xmlns:soapenv="http://schemas.xmlsoap.org/</pre>
 soap/envelope/"
   xmlns:xsd="http://www.w3.org/2001/XMLSchema"
   xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance">
 <soapenv:Body>
    <qetIDPListResponse</pre>
        soapenv:encodingStyle=
            "http://schemas.xmlsoap.org/soap/encoding/">
      <getIDPListReturn
          soapenc:arrayType=
             "ns1:IDP[2]" xsi:type="soapenc:Array"
           xmlns:ns1="urn:BeanService"
           xmlns:soapenc=
             "http://schemas.xmlsoap.org/soap/encoding">
         <getIDPListReturn href="#id0" />
         <getIDPListReturn href="#id1" />
      </getIDPListReturn>
    </getIDPListResponse>
    <multiRef id="id0" soapenc:root="0"</pre>
        soapenv:encodingStyle=
           "http://schemas.xmlsoap.org/soap/encoding/"
          xsi:type="ns2:IDP"
         xmlns:soapenc=
            "http://schemas.xmlsoap.org/soap/encoding/"
          xmlns:ns2="urn:BeanService">
       <company xsi:type="xsd:string">MegaMarket</company>
      <entityId xsi:type="xsd:string">www.megamarket.com
       </entityId>
    </multiRef>
    <multiRef id="id1" soapenc:root="0"</pre>
        soapenv:encodingStyle=
           "http://schemas.xmlsoap.org/soap/encoding/"
          xsi:type="ns3:IDP" xmlns:ns3="urn:BeanService"
         xmlns:soapenc=
            "http://schemas.xmlsoap.org/soap/encoding/">
      <company xsi:type="xsd:string">Ping</company>
      <entityId
          xsi:type="xsd:string">pingfederate3:default:entityId
       </entityId>
    </multiRef>
 </soapenv:Body>
</soapenv:Envelope>
```

OAuth Client Management Service

PingFederate includes a REST-based web service for OAuth client management.

The OAuth client management service is provided primarily for organizations with many OAuth clients, for allowing programmatic management of OAuth clients, and as an alternative to using the administrative console, the administrative API, or dynamic client registration.

The /pf-ws/rest/oauth/clients and /pf-ws/rest/oauth/clients/<clientId> REST resources are URL path extensions of the PingFederate runtime endpoint:

- https://www.example.com:9031/pf-ws/rest/oauth/clients
- https://www.example.com:9031/pf-ws/rest/oauth/clients/<clientId>





Applications must authenticate to this web service using HTTP Basic authentication and credentials validated through an instance of a Password Credential Validator (PCV). The PCV instance, in turn, must be selected in the OAuth authorization server (AS) configuration.



The administrative API can also manage OAuth clients programmatically regardless of whether the client records are managed in XML files or in a database.

Endpoint: /pf-ws/rest/oauth/clients

This resource accepts the POST, PUT, and GET methods. The POST and PUT methods described in this section require parameter name/value pairs formatted in JSON.

POST

Use the POST method to create a new client based on the parameters provided in the request. Parameters correspond to the fields on the Client window. The required MIME type is application/json.

JSON Parameters

Parameter	Description
clientId	A unique identifier the client provides to the resource server (RS) to
(Required)	identify itself. This identifier is included with every request the client makes.
enabled	Specifies whether the client is enabled. The default value is true. A valid value is either true or false.
name	A descriptive name for the client instance. This name appears when the user is prompted for authorization.
(Required)	
description	A description of what the client application does. This description appears when the user is prompted for authorization.

Parameter Description The authentication method that the client uses. clientAuthnType Set to none if your use case does not require client authentication. Note: A value other than none is required for any of the following use cases: • This client uses the client credentials grant type. See the grantTypes parameter. This client signs its ID tokens using an HMAC signing algorithm. See the idTokenSigningAlgorithm parameter. This client can access the Session Revocation API endpoint. See the grantAccessSessionRevocationApi parameter. This client sends a **secret** parameter value. Set to SECRET for HTTP Basic authentication. This authentication method requires the secret parameter. Set to CLIENT CERT for mutual SSL/TLS authentication: recommended for client applications where security policies prohibit storing passwords. This authentication method requires the clientCertIssuerDn and clientCertSubjectDn parameters. Important: If you choose mutual SSL/TLS authentication, you must configure a secondary PingFederate HTTPS port. For more information, see the property pf.secondary.https.port in the table under Configuring PingFederate properties on page 752. Set to PRIVATE KEY JWT check box if the client authenticates through the private_key_jwt client authentication method, as defined in Client Authentication in the OpenID Connect specification. For more information, see openid.net/specs/ openid-connect-core-1_0.html#ClientAuthentication. The client password or phrase. secret Required when the clientAuthnType parameter is set to

SECRET.

Description

clientCertIssuerDaThe issuer distinguished name (DN) of the client certificate.

These are certificate authority (CA) certificates imported into PingFederate on the Security # Certificate & Key Management # Trusted CAs window. Alternatively, it might be set to Trust Any to trust all the issuers found on the Trusted CAs window.

Required when the clientAuthnType parameter is set to CLIENT CERT.

clientCertSubjectThe subject DN of the client certificate.

Required when the clientAuthnType parameter is set to CLIENT CERT.

tokenEndpointAuthsthersignigad adportition that the client must use to sign the JSON web tokens (JWT) for client authentication.

> Applicable only when the clientAuthnType parameter is provided with a value of PRIVATE KEY JWT.

PingFederate accepts the following values:

- RS256 RSA using SHA-256
- RS384 RSA using SHA-384
- RS512 RSA using SHA-512
- ES256 ECDSA using P256 Curve and SHA-256
- ES384 ECDSA using P384 Curve and SHA-384
- ES512 ECDSA using P521 Curve and SHA-512
- PS256 RSASSA-PSS using SHA-256
- PS384 RSASSA-PSS using SHA-384
- PS512 RSASSA-PSS using SHA-512



Note:

RSASSA-PSS signing algorithms require either a Java 11 runtime environment or an integration with a hardware security module (HSM) and a static-key configuration for OAuth and OpenID Connect. For more information on HSM integration and static keys, see and, respectively.

If this parameter is not provided, the client can use any of the supported signing algorithms.

enforceReplayPreventsifiedd determines whether PingFederate mandates a unique signed JWT from the client for each request when the client is configured to authenticate via the private_key_jwt client authentication method, to transmit request parameters using in signed request objects, or to do both.

A valid value is either true or false.



Note:

The underlying Assertion Replay Prevention Service is clusteraware. For more information, see .

Description

requireSignedRequestarmines whether the client must transmit request parameters in a single, self-contained parameter. The parameter name is request. The value of the request parameter is a signed JWT whose claims represent the request parameters of the authorization request. The OpenID Connect specification calls this JWT a request object.

A valid value is either true or false.



Note:

If a client includes in an authorization request a request parameter, other than client id and response type, as a parameter outside of the signed request object and a claim inside of the signed request object, PingFederate always uses the claim value found inside the signed request object to process the request further.

For the client id and response type request parameters, the values outside of the signed request object must match the claim values inside of the signed request object. If the values do not match, PingFederate returns an error message to the client.

If a request parameter is found only outside of the signed request object. PingFederate ignores the request parameter and returns no error message.



Per OAuth and OpenID Connect specifications, a client must always include in an authorization request the client id, response type, and scope request parameters outside of the signed request object.

For more information about request objects, see OpenID Connect specifications.

If this parameter is not provided, a value of false is assumed.

Description

requestObjectSignInexigning almorithm that the client must use to sign its request objects for transmission of request parameters.

> Applicable only when the client might send its authorization requests using request objects.

PingFederate accepts the following values:

- RS256 RSA using SHA-256
- RS384 RSA using SHA-384
- RS512 RSA using SHA-512
- ES256 ECDSA using P256 Curve and SHA-256
- ES384 ECDSA using P384 Curve and SHA-384
- ES512 ECDSA using P521 Curve and SHA-512
- PS256 RSASSA-PSS using SHA-256
- PS384 RSASSA-PSS using SHA-384
- PS512 RSASSA-PSS using SHA-512



Note:

RSASSA-PSS signing algorithms require either a Java 11 runtime environment or an integration with a hardware security module (HSM) and a static-key configuration for OAuth and OpenID Connect. For more information on HSM integration and static keys, see and, respectively.

If this parameter is not provided, the client can use any of the supported signing algorithms.

jwksUrl, or jwks

The URL of the JSON Web Key Set (JWKS) or the actual JWKS from the client.

If the client is configured to use the private_key_jwt client authentication method, to transmit request parameters in signed request objects, or to transmit CIBA request parameters in signed request objects, only one of the previous values is required for PingFederate to verify the authenticity of the JWTs.

Either value can be defined even if the client is not configured to use JWTs for authentication or transmission of request parameters. This flexibility allows the client to transmit request parameters in signed request objects for some requests and without the use of signed request objects for some other transactions. For information on runtime processing, see .

If the client signs its JWTs using an RSASSA-PSS signing algorithm, PingFederate must be deployed to run in a Java 11 runtime environment, or be integrated with a hardware security module (HSM) to process the digital signatures. For more information on HSM integration, see .

If the client is configured to encrypt ID tokens using an asymmetric encryption algorithm, either the JWKS URL or the actual JWKS must be provided. See the ID Token Key Management Encryption Algorithm setting.

Parameter	Description
redirectUris	URIs where the OAuth AS may redirect the resource owner's user agent after authorization is obtained. The authorization code and implicit grant types require at least one redirection URI.
logoUrl	The location of the logo used on user-facing OAuth grant authorization and revocation pages. For best results with the installed HTML templates, the recommended size is 72 x 72 pixels.
allowAuthenticat	i են բարին Հրա և եր, the client can initiate an authentication API OAuth flow through the authorization endpoint without needing to handle HTTP redirections.
	When enabling this feature, consider the following:
	 redirectUris are optional, but without a redirectUri, browser-based OAuth flows will not work. This flow does not support the user-facing scope consent page, Request for Approval. So ensure that both bypassApprovalPage and restrictScopes are set to true. Also ensure that restrictedScopes is provided with an array of scope values that you allow the client to use. The client must manage the PF cookie and, if persistent authentication sessions are configured, the PF.PERSISTENT cookie.
bypassApprovalPa	glés set to true, resource-owner approval for client access is assumed, and PingFederate no longer presents to the user an authorization consent page or redirects to a trusted web application that is responsible to prompt the user for authorization for this client.
	A valid value is either true or false. If this parameter is not provided, a value of false is assumed.
	If this parameter is not provided, a value of false is assumed.

Parameter Description This setting controls whether all existing common scopes and restrictScopes scope groups, and those created in the future, or only the select few should be made available to the client. A valid value is either true or false. When set to true, PingFederate limits the client to a list of common scopes and scope groups as specified by the restrictedScopes parameter. When set to false, all existing common scopes and scope groups and those created in the future are available to the client. If this parameter is not provided, a value of false is assumed. Note: Depending on the configured dynamic scope patterns and whether they are defined as common or exclusive dynamic scopes, this setting and the restrictedScopes parameter value can impact the results of scope evaluation. The default scope, however, is always allowed for and available to all clients. For detailed information, see the **Dynamic scope evaluation and per-client** scope management section in Scopes and scope management on page 486. restrictedScopes Used in conjunction with the restrictScopes parameter value of true to limit this client to a list of common scopes or scope groups in addition to the default scope. The rest and any common scopes and scope groups created in the future become invalid for the client; that is, if the client tries to use such scope or scope group, it will receive an invalid scope error message from PingFederate. This setting controls whether any exclusive scopes and scope exclusiveScopes groups should be made available to the client. As needed, provide this parameter with a list of exclusive scopes or scope groups that are intended for the client. The rest and any exclusive scopes and scope groups created in the future become invalid for the client. In other words, if the client tries to use such scope or scope group, it will receive an invalid scope error message from PingFederate. If this parameter is not provided, no exclusive scopes or scope groups are available to the client. Note: Depending on the configured dynamic scope patterns and whether they are defined as common or exclusive dynamic scopes, this setting can impact the results of scope evaluation. The default scope, however, is always allowed for and available to all clients. For detailed information, see the Dynamic scope evaluation and per-client scope management section in Scopes and scope management on page 486.

Description

defaultAccessToke Determines the default Access Token Management (ATM) instance for this client.

validateUsingAllEApplicableAtons/ to resource server clients.

If selected, this resource server client is not required to specify the additional access token manager id or aud parameters to disambiguate the ATM instance in its token validation requests. When the resource server client does not specify the desired ATM instance, PingFederate validates the access tokens against all eligible ATM instances. This simplifies interactions with PingAccess by avoiding the need to align resource URIs between PingAccess and PingFederate.

This check box is not selected by default.

requireProofKeyFoApplicable Nathemethe client is configured to support the authorization code grant type.

A valid value is either true or false.

Determines whether the client must provide certain parameters to reduce the risk of authorization code interception attack. For more information, see the Proof Key for Code Exchange (PKCE) by OAuth Public Clients specification.

When enabled, this client must include a one-time string value through the use of the code challenge parameter in its authorization request. For more information, see . It must also submit the corresponding code verifier through the code verifier parameter in its token request when exchanging an authorization code for an access token. For more information, see.

If this parameter is not provided, the assumed value is false.

persistentGrantExpluentdeanthepeersistent Grant Max Lifetime value set globally in System # OAuth Settings # Authorization Server Settings.



Note:

This setting can be overridden per grant-mapping configuration through the use of an extended persistent grant attribute PERSISTENT GRANT LIFETIME. The PERSISTENT GRANT LIFETIME attribute is defined in System # OAuth Settings # Authorization Server Settings. When this attribute is active, you can set the lifetime of persistent grants based on the outcome of attribute mapping expressions in individual grantmapping configurations. For grant-mapping configurations that do not require this fine-grain control, you can configure them to use the default value.

persistentGrantExAminategenTapmesenting units of time for storage of persistent grants for this client.

> Required when the persistentGrantExpirationType parameter is provided with a value of OVERRIDE SERVER DEFAULT.

Description

persistentGrantExpiritation time set by the persistentGrantExpirationTime parameter.

Allowed values:

- h (hours)
- d (days)
- n (minutes)

Required when the persistentGrantExpirationType parameter is provided with a value of OVERRIDE_SERVER_DEFAULT.

persistentGrantIdDertimesattreyPersistent Grant Idle Timeout field value set globally in System # OAuth Settings # Authorization Server Settings.

Allowed values:

- SERVER DEFAULT (the default): Use the global setting.
- NONE: Grants do not expire due to inactivity.
- OVERRIDE_SERVER_DEFAULT: Use in conjunction
 with the persistentGrantIdleTimeout and
 persistentGrantIdleTimeoutUnit parameters to set the
 idle timeout window.

If an idle timeout value is configured, the idle timeout window slides when a persistent grant is updated. For more information, see .

If you configure an idle timeout value, the idle timeout window slides when a persistent grant updates. When you have an idle timeout value configured without a maximum lifetime, persistent grants remain valid until they expire due to inactivity or until the grant storage revokes or removes them. When you have an idle timeout value configured with a maximum lifetime, persistent grants remain valid until they expire due to inactivity or lifetime expiration or until the grant storage removes them.

persistentGrantIdAeTintegentepresenting the inactivity timeout value for this client.

Required when the persistentGrantIdleTimeoutType parameter is provided with a value of OVERRIDE SERVER DEFAULT.

persistentGrantIdleitsferotheTitacitwity timeout value set by the persistentGrantIdleTimeout parameter.

Allowed values:

- h (hours)
- d (days)
- n (minutes)

Required when the persistentGrantIdleTimeoutType parameter is provided with a value of OVERRIDE SERVER DEFAULT.

OpenID Connect client settings



Note:

The following parameters are only applicable when this client supports the OpenID Connect use cases.

idTokenSigningAlgothettl660N Web Signature (JWS) algorithm required for the OpenID Connect ID tokens.

Allowed values:

- none No signing algorithm
- HS256 HMAC using SHA-256
- HS384 HMAC using SHA-384
- HS512 HMAC using SHA-512
- ES256 ECDSA using P256 Curve and SHA-256
- ES384 ECDSA using P384 Curve and SHA-384
- ES512 ECDSA using P521 Curve and SHA-512
- RS256 RSA using SHA-256
- RS384 RSA using SHA-384
- RS512 RSA using SHA-512
- PS256 RSASSA-PSS using SHA-256
- PS384 RSASSA-PSS using SHA-384
- PS512 RSASSA-PSS using SHA-512



Note:

RSASSA-PSS signing algorithms require either a Java 11 runtime environment or an integration with a hardware security module (HSM) and a static-key configuration for OAuth and OpenID Connect. For more information on HSM integration and static keys, see and, respectively.



Important:

If static keys for OAuth and OpenID Connect are enabled, use either an RSA algorithm or an EC algorithm that has been configured with an active static key.

Description

idTokenEncryptionAbgadatithm used to encrypt or otherwise determine the value of the content encryption key.

Allowed values:

- dir Direct Encryption with symmetric key
- A128KW AES-128 Key Wrap
- A192KW AES-192 Key Wrap
- A256KW AES-256 Key Wrap
- A128GCMKW AES-GCM-128 key encryption
- A192GCMKW AES-GCM-192 key encryption
- A256GCMKW AES-GCM-256 key encryption
- ECDH-ES ECDH-ES
- ECDH-ES+A128KW ECDH-ES with AES-128 Key Wrap
- ECDH-ES+A192KW ECDH-ES with AES-192 Key Wrap
- ECDH-ES+A256KW ECDH-ES with AES-256 Key Wrap
- RSA-OAEP RSAES-OAEP

idTokenContentEncThetcontentgenciythion algorithm used to perform authenticated encryption on the plain text payload of the token.

> Required if an algorithm is provided through the idTokenEncryptionAlgorithm parameter.

Allowed values:

- A128CBC-HS256 Composite AES-CBC-128 HMAC-SHA-256
- A192CBC-HS384 Composite AES-CBC-192 HMAC-SHA-384
- A256CBC-HS512 Composite AES-CBC-256 HMAC-SHA-512
- AES-GCM-128 A128GCM
- AES-GCM-192 A192GCM
- AES-GCM-256 A256GCM

policyGroupId

The desired Open ID Connect policy.

grantAccessSessio 6 Retrocate do Aplow this client to access the Session Revocation API for back-channel session query and revocation.

A valid value is either true or false.

If this parameter is not provided, a value of false is assumed.



Note:

If clients are allowed to add sessions to the revocation list, you can enable the Check session revocation status option in the applicable Access Token Management instances for the token validation process to consider whether a session has been added to the revocation list. For more information, see .

grantAccessSession Management to access the Session Management API. The session management API lets clients get information about user sessions, extend sessions, and revoke sessions. For more information, see Session Management API on page 1120.

Parameter	Description
pairwiseUserType	Set to true to allow this client to use pairwise pseudonymous IDs. This parameter is set to false by default.
sectorIdentifier	Bepecify one HTTPS URL only. This parameter is applicable only if pairwiseUserType is set to true.
Note:	sions for Logout check how is selected in the System # OAuth

If the **Track User Sessions for Logout** check box is selected in the **System # OAuth Settings # Authorization Server Settings** window, you can provide two additional parameters to enable asynchronous front-channel logout for this client.

pingAccessLogoutClipath letrue, PingFederate sends logout requests through the browser to an OpenID Connect endpoint in PingAccess as part of the logout process.

A valid value is either true or false.

If this parameter is not provided, a value of false is assumed.

A list of additional endpoints at the relying parties as needed. PingFederate sends requests to these URIs through the browser as part of the logout process. The relying parties must return an image in their logout responses, otherwise PingFederate returns an error message or redirects to the InErrorResource parameter value, if specified..

Device Authorization Grant client settings

Description

deviceFlowSettingThiseield controls whether to use global device authorization grant settings defined on the System # OAuth Settings # Authorization Server Settings window.

Allowed values are SERVER_DEFAULT and OVERRIDE_SERVER_DEFAULT.

Set to OVERRIDE_SERVER_DEFAULT and configure any of the following settings:

userAuthzUrlOverride

This field controls whether PingFederate should use a different URL, perhaps for ease of use or branding purposes, when formulating the verification URLs to be included in its device authorization responses. For more information, see *Device authorization endpoint* on page 1077.

For example, if this field is configured with a value of https://www.example.org/welcome, PingFederate returns https://www.example.org/welcome and https://www.example.org/welcome?user_code=activationcode as the verification URIs.

After processing the device authorization response, which includes the verification URIs, the device presents one of them to the user. The user is expected to browse to the presented verification URI on a second device.



Important:

The target web server must redirect the browser to PingFederate at its user authorization endpoint. For more information, see *User authorization endpoint* on page 1080. Moreover, it must also preserve the user_code parameter value, if provided.

For example, if the base URL of your PingFederate server is https://www.example.com and this field is configured with a value of https://www.example.org/welcome, the target web server must redirect as follows:

- https://www.example.org/welcome to https:// www.example.com/as/user_authz.oauth2
- https://www.example.org/welcome?
 user_code=<activationcode> to https://www.example.com/as/user_authz.oauth2?user_code=<activationcode>

pendingAuthzTimeoutOverride

The lifetime of an activation code (the user_code parameter value) in seconds.

devicePollingIntervalOverride

The amount of time in seconds that the device waits between polling requests to the PingFederate token endpoint.

bypassActivationCodeConfirmationOverride

When PingFederate receives a verification request that includes an activation code (the user_code parameter value), it prompts the user to confirm the activation code.

This field controls whether PingFederate should skip this confirmation step.

Parameter Description Client-initiated backchannel authentication (CIBA) client settings cibaTokenDeliveryMbdeoken delivery method that the client supports. PingFederate supports poll and ping. Set to poll if the client can check for the authorization results periodically at the token endpoint. Set to ping if the client prefers to wait for a ping callback message from PingFederate as a signal that the authorization result is ready for pickup. If the CIBA grant type is enabled, this parameter is required; no value is assumed. cibaNotificationEntropediant's notification endpoint, to which PingFederate sends its ping call back messages. Required only if ping is the configured token delivery method. cibaPollingIntervapecifies the number of seconds that the client must wait between its attempts to check for the authorization results at the token endpoint. When PingFederate receives a token request within this time interval, it returns a slow down error message to the client. A valid value ranges from 1 to 3600. If the CIBA grant type is enabled, this parameter is required; no value is assumed. Specifies the CIBA request policy associated with the client. cibaPolicyId PingFederate uses CIBA request policies to determine various aspects of CIBA authentication request, such as the maximum lifetime of authentication requests, the validity of unsigned login hint tokens, and the mapping configuration of identity hints. Provide an existing CIBA policy If this parameter is not provided and the CIBA grant type is enabled, PingFederate associates the client with the CIBA request policy has been designated as the default CIBA request policy on the Applications # OAuth # CIBA Request Policies window. cibaUserCodeSupportates whether the client supports user code. The purpose of this code is to authorize the transmission of an authentication request to the user's authentication device. A valid value is either true or false.

If this parameter is not provided and the CIBA grant type is enabled,

When user code support is enabled, the associated CIBA request

user code support is not enabled.

policy must also be user code enabled.

Description

cibaRequireSignedRetermites whether the client must transmit request parameters in a single, self-contained parameter. The parameter name is request. The value of the request parameter is a signed JWT whose claims represent the request parameters of the authorization request. The OpenID Connect specification calls this JWT a request object.

A valid value is either true or false.

If this parameter is not provided and the CIBA grant type is enabled, CIBA signed requests are not required.

If CIBA signed requests are required, the client must also be configured with either the JWKS URL or the actual JWKS from the client.

cibaRequestObject3hersigning appoint than that the client must use to sign its request objects for transmission of request parameters.

PingFederate accepts the following values:

- RS256 RSA using SHA-256
- RS384 RSA using SHA-384
- RS512 RSA using SHA-512
- ES256 ECDSA using P256 Curve and SHA-256
- ES384 ECDSA using P384 Curve and SHA-384
- ES512 ECDSA using P521 Curve and SHA-512
- PS256 RSASSA-PSS using SHA-256
- PS384 RSASSA-PSS using SHA-384
- PS512 RSASSA-PSS using SHA-512



RSASSA-PSS signing algorithms require either a Java 11 runtime environment or an integration with a hardware security module (HSM) and a static-key configuration for OAuth and OpenID Connect. For more information on HSM integration and static keys, see and, respectively.

If this parameter is not provided and the CIBA grant type is enabled, the client can use any of the allowed signing algorithms.

Extended properties

Description

extendedParameterProvide values for extended client metadata fields.

```
{
  "extendedParams": {
    "entry": [
        "key": "ContactName",
        "value": {
          "elements": "J. Smith"
        }
      },
        "key": "ContactNumbers",
        "value": {
          "elements": [
            "555-123-4567",
            "555-987-6543"
          1
   ]
 },
}
```

This sample request provides a value for a single-valued client metadata field, ContactName, and multiple values for a multivalued client metadata field, ContactNumbers.

Extended client metadata fields are defined on the **System # Server** # **Extended Properties** window.

Sample JSON

```
"client": [
    {
      "secret":
 "L1u508MfeZYTvR03kcpa6ezysNEspFEtzxSAIEOT118AuNd2pnNqjkRdOXzfTFXc",
      "clientId": "SampleClient",
      "description": "This is a sample client.",
      "grantTypes": [
        "refresh token",
        "authorization code"
      "name": "Sample Client",
      "redirectUris": [
        "https://www.example.com/redirect1",
        "https://www.example.com/redirect2"
    }
 ]
}
```

Return codes

200 – Success

The response contains details as to why the client creation failed.

401 – Invalid Credentials

The user does not exist or is not authorized to create a client.

500 – Internal Server Error

An unknown error has occurred.

PUT

Updates client details for a specified client.



You cannot update a client ID value. You can delete the client record and create a new one with a new client ID value.

JSON Parameters

The same parameters described for *POST* apply for PUT with one addition: forceSecretChange.

Use this parameter, set to true, in conjunction with the secret parameter to change a client pass phrase.

A valid value is either true or false.

If this parameter is not provided, a value of false is assumed.



Note:

If the secret parameter is used without a forceSecretChange parameter value of true, the secret value is ignored.

Sample JSON

```
"client": [
      "secret":
 "L1u508MfeZYTvR03kcpa6ezysNEspFEtzxSAIEOT118AuNd2pnNqjkRdOXzfTFXc",
      "forceSecretChange": "true",
      "clientId": "SampleClient",
      "description": "This is a sample client.",
      "grantTypes": [
        "refresh token",
        "authorization code"
      "name": "Sample Client",
      "redirectUris": [
        "https://www.example.com/redirectOne",
        "https://www.example.com/redirectTwo"
 ]
}
```

Return codes

200 – Success

The body contains a list of updated clients.

400 - Failed To Update Client

The response contains details as to why the client could not be updated.

401 – Invalid Credentials

The user does not exist or is not authorized to update a client.

500 – Internal Server Error

An unknown error has occurred.

GET

Retrieves details for all OAuth clients.

JSON Parameters

None.

Return codes

• 200 - Success

The body contains JSON data for all clients.



The parameter refreshRolling is not returned if the AS global setting is set for a client (the default).

400 – Failed To Retrieve Clients

The response contains details as to why clients could not be retrieved.

401 – Invalid Credentials

The user does not exist or is not authorized.

500 – Internal Server Error

An unknown error has occurred.

Endpoint:/pf-ws/rest/oauth/clients/<clientId>

This resource accepts the GET and DELETE methods.

GET

Retrieves details for the specified client ID.

JSON Parameters

None.

Return codes

JSON client parameters are included.



Note:

The parameter refreshRolling is not returned if the AS global setting is set for a client, the default setting.

400 – Failed To Retrieve Client

The response contains details as to why client could not be retrieved.

401 – Invalid Credentials

The user does not exist or is not authorized.

500 – Internal Server Error

An unknown error has occurred.

DELETE

Deletes records for the specified client ID.

JSON Parameters

None.

Return codes

- 200 Success
- 400 Failed To Delete Client

The response contains details as to why client could not be deleted.

401 – Invalid Credentials

The user does not exist or is not authorized.

405 – Method Not Allowed

The client ID was not specified.

500 – Internal Server Error

An unknown error has occurred.

Logging

PingFederate records the actions performed through this endpoint in the runtime-api.log file. While the events themselves are not configurable, you can adjust the log4j2.xml configuration settings to alter the format and desired level of detail surrounding each event.

Each log entry contains information relating to the event, including:

- Time the event occurred on the PingFederate server
- Administrator username performing the action
- Authentication method
- Client IP
- HTTP method
- REST endpoint
- HTTP status code

Each of the above fields is separated by a vertical pipe (|) for ease of parsing.

OAuth Access Grant Management Service

PingFederate includes a REST-based web service to manage OAuth persistent grants administratively.

The OAuth access grant management service enables retrieval and revocation of persistent grants and their associated extended attribute names and values, if any, on a per-client or per-resource owner basis.



Important:

Client records must be stored in an external client data store.



Applications must authenticate to this web service using HTTP Basic authentication and credentials validated through an instance of a Password Credential Validator (PCV). The PCV instance, in turn, must be selected in the OAuth authorization server (AS) configuration.

Endpoints

Manage by client

/pf-ws/rest/oauth/clients/<clientId>/grants[/<grantId>]

Manage by user (based on the USER_KEY value)

/pf-ws/rest/oauth/users/<userKey>/grants[/<grantId>]

Both REST resources accept the GET, for retrieval methods, and DELETE, for revocation methods. The results are formatted in JSON.

Parameters

Parameter	Description
clientId	The client ID for which to retrieve or revoke grants.
	Required to manage grants for a specific client.
userKey	The user_key value for which to retrieve or revoke grants.
	Required to manage grants for a specific resource owner.
	Tip: The USER_KEY value varies, depending on its fulfillment based on the mapping configuration defined in the IdP Adapter Grant Mapping, Authentication Policy Contract Mapping, or Resource Owner Credentials Grant Mapping tabs.

Description
The persistent grant identifier for which to retrieve or revoke a specific
grant.
The value corresponds to the value of the id field found in the JSON array of grants returned from a previous GET / oauth/clients/ <clientid>/grants or GET /oauth/users/<userkey>/grants request.</userkey></clientid>
Note:
If this parameter is not specified, the request applies to all grants for the client or user.

Cross Site Request Forgery Protection

Both endpoints require the **x-xsrf-header** HTTP Header with any value to prevent cross site request forgery.

Example 1

Sample request

A request to retrieve all grants for client with an ID value of ac_client. Note that this client is configured with Allowed Grant Types values of Authorization Code and Refresh Token.

```
GET /pf-ws/rest/oauth/clients/ac_client/grants HTTP/1.1
Host: localhost:9031
Authorization: Basic YWRtaW46MkZlZGVyYXRl
X-XSRF-HEADER: PingFederate
```

Sample response

```
{
   "items": [
        {
            "id": "ixqWL3k9ZnPxjTaphcFLrVqwdrtvc6tV",
            "userKey": "asmith",
            "grantType": "AUTHORIZATION CODE",
            "scopes": [],
            "clientId": "ac client",
            "issued": "2019-03-15T20:00:27.343Z",
            "updated": "2019-03-15T20:00:27.343Z",
            "grantAttributes": [
                     "name": "pgeaAttrMulti",
                     "values": [
"CN=group1, OU=Resources, DC=example, DC=local",
"CN=group2, OU=Resources, DC=example, DC=local"
                },
                    "name": "pgeaAttrSingle",
                     "values": [
                         "asmith@example.local"
```

Example 2

Sample request

A request to retrieve all grants for client with an ID value of im_client. This client is configured with an Allowed Grant Types value of Implicit, and PingFederate is configured to reuse existing persistent access grants for the implicit grant type on the System # OAuth Settings # Authorization Server Settings window.

```
GET /pf-ws/rest/oauth/clients/im_client/grants HTTP/1.1
Host: localhost:9031
Authorization: Basic YWRtaW46MkZlZGVyYXRl
X-XSRF-HEADER: PingFederate
```

Sample response

For more information about persistent grants and their relationship with various grant types, see *Transient grants and persistent grants* on page 79.

Return codes

- 200 Success
- 204 Success with no content returned

Returned when revoking a persistent grant.

401 – Invalid Credentials

The user does not exist, the password is incorrect, or the user account is locked.

404 – Not Found

Returned when the requested persistent grant or client is not found.

500 – Internal Server Error

An unknown error has occurred.

Logging

PingFederate records the actions performed through this endpoint in the runtime-api.log file. While the events themselves are not configurable, you can adjust the log4j2.xml configuration settings to alter the format and desired level of detail surrounding each event.

Each log entry contains information relating to the event, including:

- Time the event occurred on the PingFederate server
- Administrator username performing the action
- Authentication method
- Client IP
- HTTP method
- REST endpoint
- HTTP status code

Each of the above fields is separated by a vertical pipe (|) for ease of parsing.

OAuth Persistent Grant Management API

This REST-based runtime API facilitates the use case where clients can assume the responsibility of grant management if the users authorize the clients to do so.

In this scenario, a client prompts the user to approve a specific scope for managing persistent grants on the user's behalf. If the user approves, the client requests an access token with such scope from PingFederate. As long as the access token remains valid, the client can retrieve and revoke persistent grants and their associated extended attribute names and values, if any, for that user.

The activation of the Persistent Grant Management API requires two settings: the scope required to perform this task and the Access Token Management instance under which the access tokens issued can be used by clients to manage persistent grants. You can find both settings from the System tab on the System # OAuth Settings # Authorization Server Settings window. Additionally, clients that are capable of managing grants on their users' behalf must be configured to use the selected Access Token Management instance in their client records.



Note:

To use this runtime API, clients must authenticate by presenting a valid access token with the required scope. The token is presented as a bearer token through the HTTP Bearer authentication scheme.



Not all OAuth use cases involve persistent grants. For more information, see Transient grants and persistent grants on page 79.

Endpoint: /pf-ws/rest/oauth/grants[/<grantId>]

This REST resource accepts the GET for retrieval methods, and DELETE for revocation methods. The results are formatted in JSON.

Parameter

Parameter	Description
grantId	The persistent grant identifier to retrieve or revoke a specific grant.
	The value corresponds to the value of the id field found in the JSON array of grants returned by a previous GET request from this endpoint.

Cross Site Request Forgery Protection

This endpoint requires the **x-xsrf-header** HTTP Header with any value to prevent cross-site request forgery.

Sample request

A request to retrieve all grants.

```
GET /pf-ws/rest/oauth/grants/ HTTP/1.1
Host: localhost:9031
Authorization: Bearer eyJhbG...IKqMfg
X-XSRF-HEADER: PingFederate
```

In this example, the client prompted the resource owner (Joe) to authorize it to manage persistent grants on his behalf. Joe agreed and approved the requested adminscope. The client then obtained an access token with the scope from PingFederate. The issued access token is a self-contained JSON web token (JWT).

This sample request illustrates a GET request from the client. The client wants to retrieve a list of grants associated with Joe by presenting the access token to the Persistent Grant Management API for authentication. The access token is truncated for readability.

Sample response

```
{
    "items": [
        {
             "id": "5a4nszZOppgo9RfRtrVXNY0Eq5ka1YZ6",
             "userKey": "joe",
"grantType": "RESOURCE_OWNER_CREDENTIALS",
             "scopes": [
                 "phone"
             "clientId": "ro client",
             "issued": "2018-12-15T00:54:30.190Z",
             "updated": "2018-12-15T00:54:30.190Z"
        },
             "id": "PTfURLoaXC97JXU6uilAORSkFQoMOLyV",
             "userKey": "joe",
"grantType": "AUTHORIZATION_CODE",
             "scopes": [
                 "openid",
                 "profile",
                 "admin"
             "clientId": "ac oic_client",
             "issued": "2018-12-15T18:29:26.018Z",
             "updated": "2018-12-15T18:29:26.018Z"
        },
             "id": "k1oFbxujlGHbfEBfVqDj0aIDllFBzqhX",
```

In this example, PingFederate returned four persistent grants associated with the resource owner from three clients.

Return codes

- 200 Success
- 204 Success with no content returned

Returned when revoking a persistent grant.

401 – Invalid Credentials

The access token is invalid (including the lack of the required scope) or missing.

404 – Not Found

Returned when the requested persistent grant is not found or a grant ID is missing.

500 – Internal Server Error

An unknown error has occurred.

Logging

PingFederate records the actions performed through this endpoint in the runtime-api.log file. While the events themselves are not configurable, you can adjust the log4j2.xml configuration settings to alter the format and desired level of detail surrounding each event.

Each log entry contains information relating to the event, including:

- Time the event occurred on the PingFederate server
- Administrator username performing the action
- Authentication method
- Client IP
- HTTP method
- REST endpoint
- HTTP status code

Each of the above fields is separated by a vertical pipe (|) for ease of parsing.

Session Management API

The runtime session management API lets OAuth clients get information about user sessions, extend sessions, and revoke sessions. Knowing the remaining lifetime of a valid session could, for example, let an application prompt the user to extend the session.

An OAuth client can take the pi.sri session identifier from its access token and send it to the session management API in an HTTP GET request. Then the API returns information in JSON format about each session associated with the pi.sri, such as:

- Last activity time
- Idle and maximum timeout time
- Authentication source that created the session



The API response body includes only sessions that were configured with the authentication sessions capability described in Configuring authentication sessions on page 279.

An OAuth client can also send the pi.sri to the session management API in an HTTP POST request to extend or revoke the sessions.

The session management API works with sessions stored in persistent storage and across clustered nodes. For this API, the runtime APIs audit log only records session revoke events.



Important:

OAuth clients must authenticate to the API using their configured client authentication method.

To configure PingFederate so that an OAuth client can use the session management API:

- Allow the client to access the session management API, as described in Configuring an OAuth client on page 515
- Configure the client's access token manager instance to include a pi.sri in access tokens, as described in Managing session validation settings on page 557

Session management API endpoints

The session management API has three endpoints, which all require the sri parameter.

The OpenID Provider configuration endpoint /.well-known/openid-configuration provides configuration information for OAuth clients to access the session management API endpoints. For more information, see OpenID Provider configuration endpoint on page 1081.

The session management API endpoints return several response codes, including:

- 200 OK: The request was successfully processed
- 400 Bad request: The format of the SRI is invalid

Endpoint /pf-ws/rest/sessionMgmt/sessions/{sri}

Use HTTP GET requests to get information about all sessions associated with the pi.sri specified by the sri parameter.

Sample request:

GET /pf-ws/rest/sessionMgmt/sessions/qzTEiEroxdzAufjYKQawm72lcBE..4RbA HTTP/1.1

```
Host: www.example.com
X-XSRF-Header: PingFederate
Authorization: Basic YWNfY2xpZW50OmdPWDh0NEQ...h3ZjI=
Cookie: PF=K60mOoBlTvWcD4frFzcKF5
```

Sample response 1:

```
"sri": "qzTEiEroxdzAufjYKQawm72lcBE..4RbA",
   "status": "HAS VALID SESSIONS",
   "lastActivityTime": "2020-06-10T17:25:00.461Z",
    "authnSessions": [ // This section can include multiple sessions
            "authnSource": {
                "sourceType": "IDP CONN",
                "id": "L07d8fse7dslShd6d 20HA8jP6",
                "entityId": "Amazon Africa A" // Only for IDP CONN
 sourceType sessions
            "creationTime": "2020-06-10T17:25:00.454Z",
            "idleTimeout": "2020-06-10T18:25:00.461Z",
            "maxTimeout": "2020-06-11T01:25:00.461Z"
        },
            "authnSource": {
                "sourceType": "ADAPTER",
                "id": "HtmlFormAdapter",
                "adapterType": "HTML Form IdP Adapter" // Only for ADAPTER
sourceType sessions
            "creationTime": "2020-06-10T17:25:00.454Z",
            "idleTimeout": "2020-06-10T18:25:00.461Z",
            "maxTimeout": "2020-06-11T01:25:00.461Z"
        }
   ]
}
```

Sample response 2:

```
{
    "sri": "qzTEiEroxdzAufjYKQawm72lcBE..4RbA",
    "status": "NO_VALID_SESSIONS",
}
```

Sample response 3:

```
{
   "sri": "qzTEiEroxdzAufjYKQawm72lcBE..4RbA",
   "status": "SESSION_REVOKED",
}
```

Endpoint /pf-ws/rest/sessionMgmt/sessions/{sri}/extend

Use HTTP POST requests to extend the idleTimeout value of all sessions associated with the pi.sri specified by the sri parameter.

Sample request:

```
POST /pf-ws/rest/sessionMgmt/sessions/qzTEiEroxdzAufjYKQawm72lcBE..4RbA/extend HTTP/1.1
Host: www.example.com
X-XSRF-Header: PingFederate
```

```
Authorization: Basic YWNfY2xpZW50OmdPWDh0NEQ...h3ZjI= Cookie: PF=K60mOoBlTvWcD4frFzcKF5
```

Sample response:

```
"sri": "gzTEiEroxdzAufjYKQawm721cBE..4RbA",
   "status": "HAS VALID SESSIONS",
    "lastActivityTime": \overline{"}2020-06-10T18:15:00.4612", // Updated
    "authnSessions": [ // This section can include multiple sessions
            "authnSource": {
                "sourceType": "IDP CONN",
                "id": "L07d8fse7dslShd6d 20HA8jP6",
                "entityId": "Amazon Africa A" // Only for IDP CONN
 sourceType sessions
            "creationTime": "2020-06-10T17:25:00.454Z",
            "idleTimeout": "2020-06-10T19:15:00.071Z", // Extended
            "maxTimeout": "2020-06-11T01:25:00.461Z"
        },
            "authnSource": {
                "sourceType": "ADAPTER",
                "id": "HtmlFormAdapter",
                "adapterType": "HTML Form IdP Adapter" // Only for ADAPTER
sourceType sessions
            "creationTime": "2020-06-10T17:25:00.454Z",
            "idleTimeout": "2020-06-10T19:15:00.071Z", // Extended
            "maxTimeout": "2020-06-11T01:25:00.461Z"
        }
   ]
}
```

Endpoint /pf-ws/rest/sessionMgmt/sessions/{sri}/revoke

Use HTTP POST requests to revoke all sessions associated with the pi.sri specified by the sri parameter.

Sample request:

```
POST /pf-ws/rest/sessionMgmt/sessions/qzTEiEroxdzAufjYKQawm72lcBE..4RbA/revoke HTTP/1.1
Host: www.example.com
X-XSRF-Header: PingFederate
Authorization: Basic YWNfY2xpZW50OmdPWDh0NEQ...h3ZjI=
Cookie: PF=K60mOoBlTvWcD4frFzcKF5
```

Sample response:

```
{
   "sri": "qzTEiEroxdzAufjYKQawm721cBE..4RbA",
   "status": "SESSION_REVOKED",
}
```

Session Revocation API endpoint

PingFederate includes a REST-based web service for back-channel session revocation.

This service enables OAuth clients to add sessions to the revocation list or to query their revocation status. This endpoint is not part of the OAuth specification. You must select the **Allow Access to Session Revocation API** check box in the configuration for the applicable clients. This endpoint is a URL path

extension of the PingFederate runtime endpoint. For example, https://www.example.com:9031/pf-ws/rest/ sessionMgmt/revokedSris.



Important:

OAuth clients must authenticate to the web service using their configured client authentication method.



If clients are allowed to add sessions to the revocation list, you can enable the Check session revocation status option in the applicable Access Token Management instances for the token validation process to consider whether a session has been added to the revocation list. For more information, see .

Endpoint: /pf-ws/rest/sessionMgmt/revokedSris

This resource accepts the *POST* and *GET* methods. It also requires the **x-xsrf-header** HTTP header with any value to prevent cross site request forgery.



The POST method described in this section requires the element name/value pair formatted in JSON.



You can find information about the Session Revocation API endpoint in the OpenID Provider Configuration endpoint metadata: ping revoked sris endpoint.

POST

A POST request adds a session to the revocation list based on its session identifier, id, in the POST data. The ID value corresponds to that of the pi.sri element in the ID token. The required Content-Type is application/json.

Sample request

A POST request to add a session with a session identifier of abc123 to the revocation list.

```
POST /pf-ws/rest/sessionMgmt/revokedSris
Host: localhost:9031
Authorization: Basic
YWNfb21jX2NsaWVudDphYmMxMjNERUZnaG1qa2xtbm9wNDU2N3JzdHV2d3h5e1pZWFdVVDg5MTBTU1
X-XSRF-HEADER: PingFederate
Content-Type: application/json
{"id": "abc123"}
```

Return codes

201 – Created

The session is added to the revocation list.

400 – Bad Request

The x-xsrf-header HTTP header is not found in the HTTP POST request.

401 – Unauthorized

The response contains details as to why the attempt failed.

415 – Unsupported Media Type

The Content-Type: application/json HTTP header is not found in the HTTP POST request.

500 – Internal Server Error

An unknown error has occurred.

GET

A GET request sends a query for the revocation status for a session with its session identifier, id, appended to the endpoint. The ID value corresponds to that of pi.sri in the ID token.

Sample request

A GET request to guery the revocation status for a session with a session identifier of abc123.

```
GET /pf-ws/rest/sessionMgmt/revokedSris/abc123
Host: localhost:9031
Authorization: Basic
YWNfb2ljX2NsaWVudDphYmMxMjNERUZnaGlqa2xtbm9wNDU2N3JzdHV2d3h5elpZWFdVVDg5MTBTUl
X-XSRF-HEADER: PingFederate
```

If PingFederate authentication sessions are enabled, querying a valid session also extends the session lifetime by the time value specified in the global **Idle Timeout** field or the idle timeout override for the authentication source associated with the session. The latter takes precedence. For externally stored authentication sessions, this operation is optimized to only send updates to the external storage when the remaining idle timeout window is less than 75%.

Alternatively, include the optional updateActivityTime query parameter and set the value to false in the request to query the status of a session without extending its lifetime.

Example

```
GET /pf-ws/rest/sessionMgmt/revokedSris/abc123?updateActivityTime=false
Host: localhost:9031
Authorization: Basic
YWNfb2ljX2NsaWVudDphYmMxMjNERUZnaGlqa2xtbm9wNDU2N3JzdHV2d3h5elpZWFdVVDg5MTBTUlFQT25tbG:
X-XSRF-HEADER: PingFederate
```

Return codes

■ 200 – OK

{"id": "abc123"} is found in the revocation list.

400 – Bad Request

The **x-xsrf-header** HTTP header is not found in the HTTP POST request.

401 – Unauthorized

The response contains details as to why the attempt failed.

404 – Not Found

{"resultId": "session_mgmt_sri_not_revoked", "message": "The SRI has not been revoked."} if the session is not found in the revocation list. If PingFederate is configured to manage authentication sessions and the request does not come with the updateActivityTime=false query parameter, the session is extended as well.

500 – Internal Server Error

An unknown error has occurred.

Logging

PingFederate records the actions performed through this endpoint in the runtime-api.log file. While the events themselves are not configurable, you can adjust the log4j2.xml configuration settings to alter the format and desired level of detail surrounding each event.

Each log entry contains information relating to the event, including:

- Time the event occurred on the PingFederate server
- Administrator username performing the action
- Authentication method
- Client IP
- HTTP method
- REST endpoint
- HTTP status code

Each of the above fields is separated by a vertical pipe (|) for ease of parsing.

PingFederate administrative API

PingFederate includes a REST-based API for administrative functions. The administrative API provides a programmatic way to make configuration changes to PingFederate as an alternative to using the administrative console.

The configuration changes that you can make through the administrative API include, but are not limited to:

- Adapters and connections
- Authentication policy contracts
- Cluster management
- Data stores and password credential validators
- Keys and certificates
- License management
- Local administrative account management
- OAuth settings
- Server settings

For a complete list, see *Accessing the API interactive documentation* on page 1132. For known limitations, see *Release notes*.

Initial setup using the administrative API

After installing PingFederate you can make four unauthenticated administrative API requests to perform the following tasks:

- 1. A GET request to /license/agreement to retrieve a URL to the license agreement.
- 2. A PUT request to /license/agreement to accept the license agreement.
- 3. A PUT request to /license to import a license file.
- 4. A POST request to /administrativeAccounts to create the first local administrative account, for native authentication.

You must assign the User Admin administrative role, USER_ADMINISTRATOR, to the first local administrative account. Other administrative roles are optional at this point. For more information, see the interactive documentation for the administrative API Accessing the API interactive documentation on page 1132.

Once the first local administrative account is created, you can make other authenticated administrative API requests to configure various components in PingFederate.

Authentication and authorization

Similar to the administrative console, access to the administrative API is protected after initial setup. The administrative API supports various authentication and authorization options. For more information, see *Configure access to the administrative API*.

Concurrent access

The administrative API supports concurrent access. When concurrent API calls are made to modify the same API resource, such as the IdP adapter instance or the SP connection, PingFederate processes the last request made.

Logging

PingFederate records actions performed through the administrative API in the admin-api.log file. Information includes the time of the event, the action performed, the authentication method, and other fields. For more information, see *Administrative API audit log*.

Configure access to the administrative API

Similar to the administrative console, access to the administrative API after initial setup might be protected by several authentication and authorization schemes.

Access to the administrative API after initial setup is protected by one of the following authentication and authorization schemes:

- Native authentication, against local administrative accounts
- LDAP authentication
- RADIUS authentication
- Mutual TLS client certificate-based authentication
- OAuth 2.0 authorization

For new installations, native authentication is the default.

For upgrades, if the authentication or authorization method of the administrative API was not previously set, such as when upgrading from PingFederate 7.3 or an earlier version, the Upgrade Utility sets the value to that of the administrative console. Otherwise, it preserves the previously set value, such as when upgrading from PingFederate 8.0 to a future release.

The authentication or authorization method for the administrative API can change at a later time to any of the choices, regardless of which authentication or authorization method is chosen for the administrative console.

Besides authentication and authorization, PingFederate also provides role-based access control, as shown in the following table. The roles assigned to the accounts affect the results of the API calls.

PingFederate User Access Control

Account type	Administrative role	Access privileges
Admin	User Admin	Create users, deactivate users, change or reset passwords, and install replacement license keys.
Admin	Admin	Configure partner connections and most system settings, except the management of local accounts and the handling of local keys and certificates.



All four administrative roles are required to access and make changes through the following services:

- The /bulk, /configArchive, and /configStore administrative API endpoints
- The Configuration Archive window, accessed from System # Server, in the administrative console
- The Connection Management configuration item on the Service Authentication window, accessed from Security # System Integration

Enabling native authentication for the administrative API

When the administrative API is protected by native authentication, access to the administrative API is restricted to the users defined in the Account Management window.

About this task

The API calls must be authenticated by valid credentials over HTTP Basic authentication; otherwise, the administrative API returns an error message. The roles assigned to the users affect the results of the API calls.

Steps

1. Verify the pf.admin.api.authentication value in <pf install>/pingfederate/bin/ run.properties is set to native. Update as needed and restart PingFederate to activate this change.



Note:

In a clustered PingFederate environment, you only need to modify run.properties on the console node.

2. Sign on to the administrative console with an account that has the User Admin role.



Important:

When the administrative console is protected by an alternative console authentication, such as certificate-based, LDAP, or RADIUS authentication, most user-management functions are handled outside the scope of the PingFederate administrative console. Therefore, the administrative console disables the functionality of the System # Server # Administrative Accounts window unless the logged-on administrator has been granted User Admin permissions.

To create or manage users in this scenario, add at least one external account to the role setting userAdmin in the configuration file for the respective authentication method. When the administrator logs on to the administrative console, the Administrative Accounts window becomes available to create or manage users for the purposes of accessing the administrative API.

For more information about the alternative console authentication and the respective configuration, see Alternative console authentication on page 780.

3. On the Administrative Accounts window, create or manage users as needed, and assign various PingFederate administrative roles as indicated by the PingFederate User Access Control table. For more information, see Configure access to the administrative API on page 1126.



When assigning roles, remember that all users defined in the Administrative Accounts window can access the administrative API and the administrative console.

Enabling LDAP authentication

When the administrative API is protected by LDAP authentication, the API calls must be authenticated by valid LDAP credentials over HTTP Basic authentication; otherwise, the administrative API returns an error message.

About this task

The LDAP authentication setup, including role assignment, is available through <pf install>/ pingfederate/bin/ldap.properties. The roles assigned to the LDAP accounts affect the results of the API calls.



Note:

When you configure LDAP authentication, PingFederate does not lock out accounts based upon the number of failed sign-on attempts. The LDAP server is responsible for preventing access and is enforced according to its password lockout settings.

Steps

- 1. Verify the pf.admin.api.authentication value in <pf install>/pingfederate/bin/ run.properties is set to LDAP. Update as needed.
- 2. In the <pf install>/pingfederate/bin/ldap.properties file, change property values as needed for your network configuration. For instructions and additional information, see the comments in the file.



Important:

Remember to assign LDAP users or designated LDAP groups, or both, to at least one of the PingFederate administrative roles, as indicated in the properties file. For information about

permissions attached to the PingFederate roles, see the PingFederate User Access Control table in Configure access to the administrative API on page 1126.



Note:

When you assign roles, remember that all LDAP accounts specified in ldap.properties can access the administrative API and the administrative console.



You can also use this configuration file in conjunction with RADIUS authentication to determine permissions dynamically with an LDAP connection.

3. Restart PingFederate.



In a clustered PingFederate environment, you only need to modify run.properties and ldap.properties on the console node.

Enabling RADIUS authentication

The RADIUS protocol provides a common approach for implementing strong authentication in a clientserver configuration.

About this task

The RADIUS authentication setup is available through configuration files in the <pf install>/ pingfederate/bin directory. The administrative API supports the protocol scenario for one-step authentication, for example, appending a one time password (OTP) after the password.

When RADIUS authentication is protecting the administrative API, the API calls must be authenticated by valid credentials over HTTP Basic authentication. Otherwise, the administrative API returns an error message. The roles assigned to the accounts affect the results of the API calls.



Note:

When you configure RADIUS authentication, PingFederate does not lock out accounts based upon the number of failed logon attempts. Instead, responsibility for preventing access is delegated to the RADIUS server and enforced according to its password lockout settings.



Note:

The NAS-IP-Address attribute is added to all Access-Request packets sent to the RADIUS server. The value is copied from the pf.engine.bind.address property in run.properties. Only IPv4 addresses are supported.

Steps

1. Verify the pf.admin.api.authentication value in <pf install>/pingfederate/bin/ run.properties is set to RADIUS. Update as needed.

2. In the <pf install>/pingfederate/bin/radius.properties file, change property values as needed for your network configuration. For instructions and additional information, see the comments in the file.



Important:

Assign RADIUS users or designated RADIUS groups, or both, to at least one of the PingFederate administrative roles as indicated in the properties file. Alternatively, you can set the use.ldap.roles property to true and use the LDAP properties file, which is also in the bin directory, to map LDAP group-based permissions to PingFederate roles. For more information about permissions attached to the PingFederate roles, see the PingFederate User Access Control table in Configure access to the administrative API on page 1126.



When assigning roles, remember that all accounts specified in radius.properties can access the administrative API and the administrative console.

3. Restart PingFederate.



In a clustered PingFederate environment, you only need to modify run.properties and radius.properties on the console node.

Enabling certificate-based authentication

When client-certificate authentication is enabled, the API calls must be authenticated by X.509 client certificates; otherwise, the administrative API returns an error message.

About this task

In addition to X.509 client certificate authentication, the corresponding root certificate authority (CA) certificates must either be contained in the Java runtime or be imported into the PingFederate's Trusted CA store. For more information, see *Manage trusted certificate authorities* on page 600.

The rest of the certificate-based authentication setup, including specifying the Issuer DN of the root CA certificates and the applicable roles of the client certificates, is available through <pf install>/ pingfederate/bin/cert auth.properties. The roles assigned to the certificates affect the results of the API calls.

Steps

1. Sign on to the administrative console with an account that has the role Crypto Admin.

2. Ensure the client-certificate's root CA and any intermediate CA certificates are contained in the trusted store, either for the Java runtime, or PingFederate, or both.



Note:

To import a certificate, click Trusted CAs in the Certificate Management section under Server Configuration.



Click the Serial number and copy the Issuer distinguished name (DN) to use in a couple steps later.

- 3. Verify the pf.admin.api.authentication value in <pf install>/pingfederate/bin/ run.properties is set to cert. Update as needed.
- 4. In the <pf install>/pingfederate/bin/cert auth.properties file, enter the Issuer DN for the client certificate as a value for the property: rootca.issuer.<x>, where <x> is a sequential number starting at 1. For more information, see the properties file.



Important:

The configuration values are case-sensitive.

If you copied the Issuer DN a couple steps earlier, paste this value.

- 5. Repeat the previous step for any additional CAs as needed.
- 6. Enter the certificate's Subject DN for the applicable PingFederate permission roles, as described in the properties file. For information about permissions attached to the PingFederate roles, see the PingFederate User Access Control table in Configure access to the administrative API on page 1126.



Important:

The configuration values are case-sensitive.



Note:

When assigning roles, keep in mind that all client certificates specified in cert auth.properties can be used to access the administrative API and the administrative console.

- 7. Repeat the previous step for all client certificates as needed.
- 8. Restart PingFederate.



Note:

In a clustered PingFederate environment, you only need to modify run.properties and cert auth.properties on the console node.

Enabling OAuth 2.0 authorization

PingFederate clients can gain access to the administrative API endpoint by providing an OAuth 2.0 access token. The <pf install>/pingfederate/bin/oauth2.properties file contains settings that allow you to configure information required to interact with the authorization server as a client.

Steps

- 1. Set the pf.admin.api.authentication value in <pf install>/pingfederate/bin/ run.properties to OAuth2.
- 2. In the <pf install>/pingfederate/bin/oauth2.properties file, change property values as needed. For instructions and additional information, see the comments in the file.



Important:

Remember to assign at least one of the PingFederate administrative roles, as indicated in the properties file. For information about permissions attached to the PingFederate roles, see the PingFederate User Access Control table in Configure access to the administrative API on page 1126.

3. Restart PingFederate.



Note:

In a clustered PingFederate environment, you only need to modify run.properties and oauth2.properties on the console node.

Accessing the API interactive documentation

PingFederate ships with interactive documentation for both developers and non-developers to explore the API endpoints, view documentation for the API, and experiment with API calls.

About this task

In general, you can make API calls from an interactive user interface, custom applications, or from command line tools such as cURL. The endpoint is only available at the administrative port, as defined by the pf.admin.https.port property in <pf install>/pingfederate/bin/run.properties.



Important:

For enhanced API security, you must include x-xsrf-Header: PingFederate in all requests and use the application/json content type for PUT and POST requests.

To access the administrative API documentation, follow these steps:

Steps

- 1. Start PingFederate.
- 2. Start a web browser.
- 3. Browse to the following URL: https://<pf_host>:9999/pf-admin-api/api-docs/



Note:

<pf_host> is the network address of your PingFederate server. It can be an IP address, a host name, or a fully qualified domain name. It must be reachable from your computer.

9999 is the default value of the pf.admin.https.port property in the run.properties file.

- 4. To test an administrative API, follow these steps:
 - a. Select a section of the administrative API you would like to explore; for example, /dataStores.
 - b. Expand the method you want to use; for example, GET /dataStores.
 - c. Enter required parameters, if any. For more information, see Operation Models underneath the selected API endpoint.
 - d. Click **Try it out**.



Note:

You might be prompted to sign on using administrative credentials over HTTP Basic authentication. The role assigned to the respective administrative accounts affects the access to the requested API.

Result:

If the request completes successfully, the administrative API returns the Request URL, the Response Body, the Response Code, and the Response Headers.

Application endpoints

Application endpoints provide a means, through standard HTTP, by which external applications can communicate with the PingFederate server.

The single sign-on (SSO) and single log-out (SLO) endpoints for an identity provider (IdP) and a service provider (SP) include optional parameters which you can use to specify error pages that users see in the event of an SSO or SLO failure. By default, PingFederate provides templates for these and other errors or conditions. For more information, see Customizable user-facing pages on page 798.

SP endpoints also include those available for system for cross-domain identity management (SCIM) inbound provisioning. For more information, see *Provisioning for SPs* on page 98.

For either SP or IdP servers, PingFederate provides a maintenance endpoint for administrators to verify that the server is running. Endpoints applicable to both server roles include those needed for adapterto-adapter mapping and retrieval of WS-Trust metadata. For more information, see Adapter-to-adapter mappings on page 371 and WSC and WSP support on page 73.

PingFederate provides a favorite icon for all application endpoints. For more information, see *Customizing* the favicon for application and protocol endpoints on page 836.

IdP endpoints

The following sections describe PingFederate identity provider (IdP) endpoints, including the case-sensitive query parameters that each accepts or requires. These endpoints accept either the HTTP GET or POST methods.

Begin each URL with the fully-qualified server name and port number of your PingFederate IdP server, for example, https://www.example.com:9031/idp/startSSO.ping.



Important:

When using the parameters TargetResource or TARGET with their own query parameters included, the parameter value must be URL-encoded. Any other parameters that contain restricted characters, such as many SAML URNs, also must be URL-encoded. For information about URL encoding, see third party resources such as HTML URL-encoding Reference. Parameters are case-sensitive.

You can customize and localize user-facing templates.

This is the path used to initiate an unsolicited IdP-initiated single sign-on (SSO) transaction during which a SAML response containing an assertion is sent to a service provider (SP). Typically, a systems integrator or developer creates one or more links to this endpoint in the IdP application or portal to allow users to initiate SSO to various SPs.

For information about allowing applications to retrieve configuration data from the PingFederate server over SOAP, see *Web service interfaces and APIs* on page 1086.

The following table shows the HTTP parameters for this endpoint.

Parameter	Description
PartnerSpId Of PARTNER	The federation ID of the SP to whom the SAML response containing an assertion should be issued. This ID value is case-sensitive.
	One of these parameters is required unless the federation ID can be derived from TargetResource or TARGET.
TargetResource Of TARGET (optional)	For SAML 2.0, the value of either parameter is passed to the SP as the RelayState element of a SAML response message. This is the PingFederate implementation of the SAML 2.0 indicator for a desired resource at the SP during IdP-initiated SSO.
	For SAML 1.x, the value is sent to the SP as a parameter named TARGET.
	The parameter value must be URL-encoded.
InErrorResource (optional)	Indicates where the user is redirected after an unsuccessful SSO. If this parameter is not included in the request, PingFederate redirects the user to the SSO error landing page hosted within PingFederate.
Binding (optional)	Indicates the binding to be used; allowed values are URIs defined in the SAML specifications. For example, the SAML 2.0 applicable URIs are:
(optional)	urn:oasis:names:tc:SAML:2.0:bindings:HTTP-Artifacturn:oasis:names:tc:SAML:2.0:bindings:HTTP-Post
	When the parameter is not used, the default ACS URL configured for the SP-partner connection is used unless an ACS index is specified using the ACSIdx parameter.
ACSIdx (optional - SAML 2.0)	Specifies the index number of partner's ACS. For more information, see Setting Assertion Consumer Service URLs (SAML) on page 430. Takes precedence over the Binding parameter if both are specified. If neither the binding nor index is specified in the call, the default ACS is used.
IdpAdapterId (optional)	Allows an application to call out what IdP adapter to use for authentication in a configuration with multiple IdP adapters.
	Note: This parameter might be overridden by policy based on authentication policies. For example, a CIDR Authentication Selector instance could enforce the use of a given adapter instance based on whether a user is on or off the network. For more information, see <i>Authentication policies</i> on page 220.

/idp/startSLO.ping

This is the path used to start an IdP-initiated SLO (under SAML 2.0) or an OpenID Connect logout. For more information see *Asynchronous Front-Channel Logout* on page 593. Typically, a systems integrator or developer creates one or more links to this endpoint in the protected resources of their IdP application or portal to allow users to end their sessions at various SPs. This endpoint uses the local PingFederate session to determine which SPs have been issued an SSO assertion and sends them a SAML logout request.

PingFederate sends SLO requests in the following sequence, which prioritizes synchronous logouts:

- 1. IdP adapters
- 2. SAML IdP partners
- 3. SP adapters
- 4. SAML SP partners
- 5. WS-Fed and OIDC partners in parallel

The **LogoutType** parameter lets you customize the SLO process. For example, you can change the SLO process so that it does not prioritize synchronous logouts.

To start with asynchronous logouts, you could chain the logouts. Start with /idp/startSLO.ping? LogoutType=AsyncOnly. Set the TargetResource to /idp/startSLO.ping? LogoutType=SyncOnly. In this case, you would also need to add the /idp/startSLO.ping endpoint as an allowed redirect for SLO on the Redirect Validation window.

To perform asynchronous and synchronous logouts in parallel, you could create a page with two iframes. One iframe points to /idp/startSLO.ping?LogoutType=AsyncOnly and the other points to /idp/startSLO.ping?LogoutType=SyncOnly.

The following table describes the HTTP parameters for this endpoint.

/idp/writecdc.ping

This endpoint is used for SAML 2.0 IdP Discovery. This is the path used when the IdP wants to write to the common domain cookie (CDC) held within the user's browser. The information written to the cookie indicates from which IdP this user has authenticated.

The following table shows the one HTTP query parameter for this endpoint.

Parameter	Description
TargetResource	Indicates where the user is redirected after successful IdP Discovery. If this
(optional)	parameter is not included in the request, PingFederate redirects the user to the referrer in the HTTP header. If there is no TargetResource or referrer, the call to this endpoint will fail.
	The parameter value must be URL-encoded.

This endpoint is used for profile management. When profile management is enabled for customer identities, authenticated users can review and modify the local identity fields that have been configured to be shown on the profile management page, connect or disconnect third-party identity providers, also known as Social Connections to the end users on the profile management page, and delete their local identities if the option to do so has been enabled. Each local identity profile has its own profile management URL.

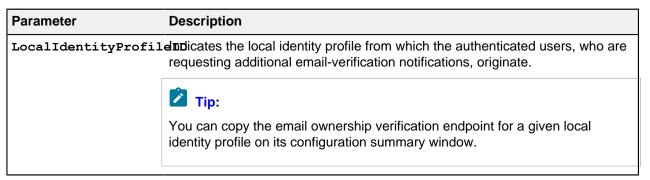
The following table shows the one HTTP query parameter for this endpoint.

Parameter LocalIdentityProfileInDicates which profile management page that PingFederate should serve to the authenticated users based on the ID of the local identity profile. Tip: You can copy the profile management URL for a given local identity profile on its configuration summary window.

/pf/id/verification.ping

This endpoint is used for email ownership verification. When email ownership verification is enabled for customer identities, authenticated users can request additional email-verification notifications by accessing this endpoint.

The following table shows the one HTTP query parameter for this endpoint.



/ext/pwdchange/Identify

The Change Password endpoint allows users to change their password through an HTML Form Adapter instance without submitting SSO requests. This endpoint requires one parameter, <code>AdapterId</code>; the parameter value is the identifier of the HTML Form Adapter instance that has been configured with such capability. For example, if the fully-qualified name of your PingFederate environment and the adapter ID are www.example.com and HTMLFormSimplePCV respectively, the resulting URL is https://www.example.com/ext/pwdchange/Identify?AdapterId=HTMLFormSimplePCV.

The following table shows the one additional HTTP guery parameter for this endpoint.

Parameter	Description
TargetResource	Indicates the desired destination after users have successfully change their network password.
	Note:
	When target resource validation is enabled for this endpoint, as indicated by the SLO and Other check box on the Security # System Integration # Redirect Validation # Local Redirect Validation tab, PingFederate honors the URL only if the parameter value satisfies the configured requirement on the Local Redirect Validation tab. If the validation fails, PingFederate displays the default success or error message.
	For more information, see <i>Configuring redirect validation</i> on page 631.

/ext/pwdreset/Identify

The Account Recovery endpoint allows users to reset their password, unlock their account, or recover their username through an HTML Form Adapter instance without submitting SSO requests. This endpoint requires one parameter, <code>AdapterId</code>; the parameter value is the identifier of the HTML Form Adapter instance that has been configured with such capabilities. For example, if the fully-qualified name of your PingFederate environment and the adapter ID is www.example.com and HTMLFormSimplePCV respectively, the resulting URL is https://www.example.com/ext/pwdreset/Identify? AdapterId=HTMLFormSimplePCV.

The following table shows the one additional HTTP query parameter for this endpoint.

Parameter	Description
TargetResource	Indicates the desired destination after users have successfully reset their network password, unlock their account, or recover their username.
	Note:
	When target resource validation is enabled for this endpoint, as indicated by the SLO and Other check box on the Security # System Integration # Redirect Validation # Local Redirect Validation tab, PingFederate honors the URL only if the parameter value satisfies the configured requirement on the Local Redirect Validation tab. If the validation fails, PingFederate displays the default success or error message.
	For more information, see <i>Configuring redirect validation</i> on page 631.

SP endpoints

PingFederate provides configuration options for a variety of endpoints.

The following sections describe the PingFederate service provider (SP) endpoints for SP services and system for cross-domain identity management (SCIM) inbound provisioning.

The following sections describe PingFederate service provider (SP) endpoints, including the guery parameters that each accepts or requires. These endpoints accept either the HTTP GET or POST methods.

Begin each URL with the fully-qualified server name and port number of your PingFederate SP server; for example, https://www.example.com:9031/sp/startSSO.ping.



Important:

When using the parameters TargetResource or TARGET with their own query parameters included, the parameter value must be URL-encoded. Any other parameters that contain restricted characters, such as many SAML URNs, also must be URL-encoded. For information about URL encoding, see third party resources such as HTML URL-encoding Reference. Parameters are case-sensitive.

/sp/startSSO.ping

This is the path used to initiate SP-initiated single sign-on (SSO). In this scenario, the SP issues an SSO request to the identity provider (IdP) asking for an SSO authentication response. Typically, a systems integrator or developer creates links to this endpoint in SP applications to allow users to access various protected resources through SSO using the IdP as an authentication authority.

For information about allowing applications to retrieve configuration data from the PingFederate server over SOAP, see Web service interfaces and APIs on page 1086.

The following table shows the HTTP parameters for this endpoint.



Some parameters described below can have multiple values. Specify these values by using multiple independent query string parameters of the same name.

Parameter	Description
PartnerIdpId	The federation ID of the IdP that authenticates the user and issues an assertion. This ID is case-sensitive.
	Required if more than one IdP connection is configured and Domain is not used, and SP authentication policies are turned off.
	Not required if SP authentication policies are turned on.
SpSessionAuthnAdapt	eElimetexplicit SP adapter instance ID indicating the adapter to use to create an authenticated session or security context.
	Optional if SP authentication policies are turned off.
	Required if SP authentication policies are turned on unless the PingFederate SP server can determine the applicable SP adapter instance based on the target URL mapping configuration and the TargetResource or TARGET value at runtime.

Parameter	Description
TargetResource Of TARGET	This parameter indicates the target applications where a successful SSO redirects the end-user.
	The parameter value must be URL-encoded.
	When this parameter is not provided in the URL, you can specify a default target resource in the administrative console, either for all IdP connections, for individual connections, or both. For more information, see <i>Configuring default URLs</i> on page 656 and <i>Configuring default target URLs</i> on page 694.
InErrorResource	This parameter indicates where an unsuccessful SSO redirects the end-user.
(optional)	If this parameter is not included in the request, PingFederate redirects the user to the single log-out (SLO) error landing page hosted within PingFederate. For more information, see <i>Customizable user-facing pages</i> on page 798.
Binding	Indicates the binding to use; allowed values are URIs defined in the SAML
(optional)	specifications. For example, the SAML 2.0 applicable URIs are
	urn:oasis:names:tc:SAML:2.0:bindings:HTTP-Artifact urn:oasis:names:tc:SAML:2.0:bindings:HTTP-POST urn:oasis:names:tc:SAML:2.0:bindings:HTTP-Redirect
	When the parameter is not used for SAML 2.0, the first SSO Service URL configured for the IdP-partner connection is used. For more information, see <i>Specifying SSO service URLs (SAML)</i> on page 685.
AllowCreate (optional - SAML 2.0)	Controls the value of the AllowCreate attribute of the NameIDPolicy element in the AuthnRequest. The default is true.
AuthenticatingIdpId (optional - SAML 2.0)	This parameter indicates the preferred IdP for authenticating the user through an IdP proxy, such as PingOne for Enterprise. The parameter specifies the value of the ProviderID attribute in the Scoping/IDPList/IDPEntry element in the AuthnRequest. For more information, see section 3.4.1.3.1 of the OASIS SAML document saml-core-2.0-os.pdf.
	You can specify multiple values to build a preferred list.
ForceAuthn (optional - SAML 2.0 or	For SAML 2.0, this parameter controls the attribute of the same name in the AuthnRequest.
OpenID Connect)	For OpenID Connect, a value of true sets the prompt parameter in the authentication request to login . For more information about the authentication request and its parameter, see the <i>OpenID Connect specification</i> .
	The default is false.
IsPassive (optional - SAML 2.0 or	For SAML 2.0, this parameter controls the attribute of the same name in the AuthnRequest.
OpenID Connect)	For OpenID Connect, a value of true sets the prompt parameter in the authentication request to none.
	The default is false.
RequestedACSIdx (optional - SAML 2.0)	The index number of your site's Assertion Consumer Service, where you want the assertion to be sent.

Parameter	Description	
RequestedAcsUrl	The URL of your site's Assertion Consumer Service, where you want the assertion to be sent.	
(optional - SAML 2.0)	assertion to be sent.	
RequestedAuthnCtx	For SAML 2.0, this parameter indicates the requested authentication context of the assertion; allowed values include URIs defined in the SAML	
(optional - SAML 2.0 or OpenID Connect)	specifications. For more information, see the OASIS SAML document saml-authn-context-2.0-os.pdf.	
	For OpenID Connect, the specified value becomes the acr_values parameter value in the authentication request.	
	You can specify multiple values to build a preferred list.	
RequestedAuthnDeclI	Refin alternative to RequestedAuthnCtx above, indicating the requested	
(optional - SAML 2.0)	authentication context of the assertion by declaring any URI reference. For more information see section 2.7.2.2 of the OASIS SAML document <i>saml-core-2.0-os.pdf</i> .	
	You can specify multiple values to build a preferred list.	
RequestedBinding	Indicates the binding requested for the response containing the assertion;	
(optional - SAML 2.0)	allowed values are URIs defined in the SAML specifications.	
RequestedFormat	Specifies the value for the Format attribute in the NameIDPolicy element	
(optional - SAML 2.0)	of the AuthnRequest. If not specified, the AuthnRequest does not include the attribute.	
RequestedSPNameQual	RequestedSPNameQualifdisates that the IdP should return the given name qualifier as part of the	
(optional - SAML 2.0)	assertion used primarily to identify SP affiliations. For more information, see SP affiliations on page 471.	
vsid	Specify the virtual server ID.	
(optional)	When absent, PingFederate uses the default virtual server ID (if specified) for the connection or the SAML federation ID defined in Server Settings . For more information, see <i>Identifying the partner</i> on page 668 and <i>Specifying federation information</i> on page 838.	

If SpSessionAuthnAdapterId specifies an adapter, then that adapter is used to create an authenticated session for SP-initiated SSO. If there is no SpSessionAuthnAdapterId, the ultimate destination of the user after SSO, either the TargetResource or the default SSO success URL, is used along with the mappings defined in the administrative console on the Map URLs to Adapter Instances window. For more information, see Configuring target URL mapping on page 652.

Note that adapter selection for SP-initiated SSO is similar to that for IdP-initiated SSO except that, because the adapter ID depends on the SAML deployment, PingFederate cannot expect it from an IdP. Therefore, it uses only the URL mapping for adapter selection for SSO.

/sp/startSLO.ping

This is the path used to initiate SP-initiated SLO. Typically, a systems integrator or developer creates one or more links to this endpoint in the protected resources of their SP application, which allows users to end a session by sending a logout request to the IdP that authenticated the session.

Note that the IdP might send additional logout request messages to other SPs when it receives a logout request from a PingFederate server acting as an SP.

The following table shows the HTTP parameters for this endpoint.

Parameter	Description
TargetResource (optional)	Indicates where a successful SLO redirects the user. If the request does not include this parameter, PingFederate uses the URL for a successful SLO as a default, as entered on the SP Default URLs window.
	Note that the parameter value must be URL-encoded.
Binding (optional - SAML 2.0)	Indicates the binding to use; allowed values are URIs defined in the SAML specifications. The SAML 2.0 applicable URIs are
	urn:oasis:names:tc:SAML:2.0:bindings:HTTP-Artifact urn:oasis:names:tc:SAML:2.0:bindings:HTTP-POST urn:oasis:names:tc:SAML:2.0:bindings:HTTP-Redirect urn:oasis:names:tc:SAML:2.0:bindings:SOAP
	When the parameter is not used, the first SLO Service URL configured for the IdP-partner connection is used. For more information, see <i>Specifying SLO service URLs (SAML 2.0)</i> on page 434.
InErrorResource (optional)	Indicates where an unsuccessful SLO redirects the user. If the request does not include this parameter, PingFederate redirects the user to the SLO error landing page hosted within PingFederate. For more information, see <i>Customizable user-facing pages</i> on page 798.

/sp/defederate.ping

This path terminates an account link created during SSO. Account linking provides a means for subject identification on the SP side. On the SP side, only users create and terminate links. The link contains the name identifier from the IdP, the IdP's federation ID, the adapter instance ID, and the local user identifier.

There are no HTTP parameters for this endpoint.

You can unlink a user session only if it was established during SSO using an existing account link on the SP side. If more than one SP session was established through account linking on the same PingFederate session, this endpoint will terminate each of those links. A local logout is also performed for any link that is terminated.

/sp/cdcstartSSO.ping

This endpoint is used for IdP-Discovery implementations. For more information, see Standard IdP Discovery on page 44. This endpoint is similar to /sp/startSSO.ping and accepts the same parameters, with the exception of PartnerIdpId and vsid. Instead of this parameter, the server attempts to use the common domain cookie to determine the IdP.

/sp/startAttributeQuery.ping

This endpoint initiates an Attribute Query with a SAML 2.0 IdP. For more information, see Attribute Query and XASP on page 43.

The following table shows the HTTP parameters for this endpoint.



Some parameters described below can have multiple values. Specify these values by using multiple independent query string parameters of the same name.

Uniquely identifies the user to the IdP. When user authenticates with an X.509 certificate, this is the Subject DN, which must be URL-encoded.
The IssuerDN from the user's X.509 certificate, when X.509 attribute sharing profile (XASP) is used, which uniquely identifies the entity that issued the
user's certificate. The parameter must be URL-encoded.
Note:
When specified this parameter overrides the Subject parameter.
Used to identify the specific IdP partner to which the attribute query should be sent. Without this parameter, the Subject and Issuer are used to determine the correct IdP.
Note:
For XASP, this parameter overrides both the Subject and Issuer parameters.
Identifies the name-identifier format of the Subject query parameter. If
included, the value must be one of the SAML 2.0 Name Identifier Format URIs For more information, see section 8.3 of the SAML specifications.
Note:
For XASP, this parameter must be set to urn:oasis:names:tc:SAML:1.1:nameid-format:X509SubjectName
If not specified, the parameter defaults to
urn:oasis:names:tc:SAML:1.1:nameid-format:unspecified.
The parameter must be URL-encoded.
The unique identifier of the initiating application.
Used to authenticate the initiating application. Both the Appld and SharedSecret values must match those defined on the Security # System Integration # Service Authentication window.
Important:
To avoid recording this parameter in web server logs, only pass it in the message body using the HTTP POST method.
A name of a user attribute requested from the IdP. For each desired user attribute, include this parameter. If this parameter is not present, then the IdP
returns all allowable user attributes.

Parameter	Description
vsid	Specify the virtual server ID.
(optional)	When absent, PingFederate uses the default virtual server ID, if specified, for the connection or the SAML federation ID defined in Server Settings . For more information, see <i>Identifying the partner</i> on page 668 and <i>Specifying federation information</i> on page 838.

SCIM inbound provisioning endpoints

PingFederate supports system for cross-domain identity management (SCIM) inbound provisioning and provides four endpoints.

The four endpoints are:

- /pf-scim/v1/Users
- /pf-scim/v1/Groups
- /pf-scim/v1/Schemas
- /pf-scim/v1/ServiceProviderConfigs

These endpoints are defined in the following SCIM 1.1 specifications:

- SCIM Core Schema
- SCIM Specification

Begin each endpoint with the fully-qualified server name and port number of your PingFederate server, for example: https://pingidentity.com:9031/pf-scim/v1/Users.

/pf-scim/v1/Users

The users endpoint is where client applications make HTTP requests to create, retrieve, update, and delete or deactivate users. This REST-based endpoint accepts POST, GET, PUT, and DELETE methods, as described in the following table.



Note:

HTTP requests must be made using either Basic or client-certificate application authentication. JSON is currently the only supported format for the HTTP message body.

HTTP method	Description
POST	/pf-scim/v1/Users
	 Sends user attributes in JSON format—defined in the SCIM Core Schema—to create a new user. A successful response is indicated by an HTTP 201 status code and a message body containing the user record that has been added to the target datastore. The user ID is set as the id attribute in the JSON response, and the full URL to reference the user is in the HTTP response Location header.
	For an existing user, you can also use the POST method to either update or delete or disable a user record by appending the user ID to the path in the format of /pf-scim/v1/Users/user_id and setting the request header x-HTTP-Method-Override value to PUT or DELETE, respectively. For more information, see the PUT and DELETE method descriptions at the end of this topic.

HTTP method Description **GET** /pf-scim/v1/Users

- Retrieves all user records.
- A successful response is indicated by an HTTP 200 status code and a list of all users and their attributes.

/pf-scim/v1/Users/user id

- Retrieves the user record of a specific user.
- A successful response is indicated by an HTTP 200 status code and the requested user record.

/pf-scim/v1/Users?attributes=attribute

- Retrieves the specific attribute from all users.
- A successful response is indicated by an HTTP 200 status code and a list of the desired attribute from all users.



Note:

For more information, see 3.2.2 List/Query Resources in SCIM Specification.

/pf-scim/v1/Users?filter=filter

- Retrieves resources based on the filter.
- A successful response is indicated by an HTTP 200 status code and a list of resources matching the filter.



Note:

For more information, see 3.2.2.1 Filtering in SCIM Specification.

/pf-scim/v1/Users?sortBy=attribute&sortOrder=ascending| descending

- Retrieves all user records and sorts them based on a specific attribute in ascending or descending order.
- A successful response is indicated by an HTTP 200 status code and a sorted result set. Depending on the implementation of the target datastore, the target datastore might not return the user records that do not contain a value for that specific attribute as indicated by the sortBy parameter in the request.



Note:

For more information, see 3.2.2.2 Sorting in SCIM Specification.

/pf-scim/v1/Users?startIndex=x[&count=y]

- Retrieves the user records starting with a specific index number, a positive integer x. If the optional count parameter is included, with a positive integer y, the endpoint limits the number of user records in the result set.
- A successful response is indicated by an HTTP 200 status code and a limited set of user records.



Note:

For more information, see 3.2.2.3 Pagination in SCIM Specification.

HTTP method	Description
PUT	/pf-scim/v1/Users/user_id
	 Updates user attributes for the specified user, using JSON in the body of the HTTP request. Attributes not included in the request are set to a default value in the datastore. A successful PUT operation returns an HTTP 200 status code and the entire updated user record within the response body.
DELETE	/pf-scim/v1/Users/user_id
	 Deletes or disables the user record for the specified user. Note that whether a user is deleted or disabled is determined by the selection of the SCIM DELETE message behavior option on the Delete/Disable Users tab in the applicable identity provider (IdP) connection. A successful response is indicated by an HTTP 200 status code.



For a list of HTTP error codes that might be returned, see 3.9 HTTP Response Codes in SCIM Specification.

/pf-scim/v1/Groups

The groups endpoint is where client applications make HTTP requests to create, retrieve, update, and delete groups.



Inbound provisioning for groups is a per-connection, optional feature. To enable group provisioning, select the User and Group Support option on the Connection Type tab when configuring the applicable IdP connection.

This REST-based endpoint accepts POST, GET, PUT, and DELETE methods, as described in the following table.



Note:

HTTP requests must be made using either Basic or client-certificate application authentication. JSON is currently the only supported format for the HTTP message body.

HTTP method	Description
POST	/pf-scim/v1/Groups
	 Sends group attributes in JSON format—defined in the SCIM Core Schema—to create a new group. A successful response is indicated by an HTTP 201 status code and a message body containing the group record that has been added to the target datastore. The group ID is set as the id attribute in the JSON response, and the full URL to reference the group is in the HTTP response Location header.
	For an existing group, you can also use the POST method to either update or delete the group by appending the group ID to the path, in the format of \pf-scim/v1/Groups/group_id, and setting the request header x-HTTP-Method-Override value to PUT or DELETE, respectively. For more information, see the PUT and DELETE method descriptions at the end of this topic.

HTTP method **Description GET** /pf-scim/v1/Groups Retrieves all group records. A successful response is indicated by an HTTP 200 status code and a list of all groups and their attributes. /pf-scim/v1/Groups/group id Retrieve the group record of a specific group. A successful response is indicated by an HTTP 200 status code and the requested group record. /pf-scim/v1/Groups?attributes=attribute

- Retrieves the specific attribute from all groups.
- A successful response is indicated by an HTTP 200 status code and a list of the desired attribute from all groups.



Note:

For more information, see 3.2.2 List/Query Resources in SCIM Specification.

/pf-scim/v1/Groups?filter=filter

- Retrieves resources based on the filter.
- A successful response is indicated by an HTTP 200 status code and a list of resources matching the filter.



Note:

For more information, see 3.2.2.1 Filtering in SCIM Specification.

/pf-scim/v1/Groups?sortBy=attribute&sortOrder=ascending| descending

- Retrieves all group records and sorts them based on a specific attribute in ascending or descending order.
- A successful response is indicated by an HTTP 200 status code and a sorted result set. Depending on the implementation of the target datastore, the target datastore might not return the group records that do not contain a value for that specific attribute as indicated by the sortBy parameter in the request.



Note:

For more information, see 3.2.2.2 Sorting in SCIM Specification.

/pf-scim/v1/Groups?startIndex=x[&count=y]

- Retrieves the group records starting with a specific index number, a positive integer x. If the optional count parameter is included, with a positive integer v, the endpoint limits the number of user records in the result set.
- A successful response is indicated by an HTTP 200 status code and a limited set of group records.



Note:

For more information, see 3.2.2.3 Pagination in SCIM Specification.



For a list of HTTP error codes that might be returned, see 3.9 HTTP Response Codes in SCIM Specification.

/pf-scim/v1/Schemas

The schemas endpoint is where a client can retrieve a resource's schema. This REST-based endpoint accepts GET method as described in the following table.



HTTP requests must be made using either Basic or client-certificate application authentication. JSON is currently the only supported format for the HTTP message body.

HTTP method	Description
GET	Retrieves the resource's schema for an IdP connection based on the authentication information.
	A successful response is indicated by an HTTP 200 status code and the results in the message body.

Example

```
$ curl -u basicUser 'https://localhost:9031/pf-scim/v1/Schemas' | python -m
json.tool
    "attributes": [
            "caseExact": false,
            "description": "Unique identifier for the SCIM resource as
defined by the Service Provider. Each representation of the resource MUST
 include a non-empty id value. This identifier MUST be unique across the
Service Provider's entire set of resources. It MUST be a stable, non-
reassignable identifier that does not change when the same resource is
 returned in subsequent requests. The value of the id attribute is always
 issued by the Service Provider and MUST never be specified by the Service
Consumer. REQUIRED.",
            "multiValued": false,
            "name": "id",
            "readOnly": true,
```

```
"required": true,
        "schema": "urn:scim:schemas:core:1.0",
        "type": "string"
    },
"description": "Core User",
"endpoint": "/Users",
"id": "urn:scim:schemas:core:1.0:User",
"name": "User",
"schema": "urn:scim:schemas:core:1.0"
```

/pf-scim/v1/ServiceProviderConfigs

This service provider (SP) configuration endpoint is where developers can retrieve detailed information on the PingFederate SCIM 1.1 implementation. When you enable inbound provisioning for an SP PingFederate server, an HTTP GET request to this endpoint returns a JSON response outlining SCIM 1.1 compliance details.



The /pf-scim/v1/ServiceProviderConfigs endpoint does not require authentication. JSON is currently the only supported format for the HTTP message body.

Example

```
$ curl https://localhost:9031/pf-scim/v1/ServiceProviderConfigs
  "schemas": ["urn:scim:schemas:core:1.0"],
  "patch": {
   "supported":false
  "bulk": {
    "supported":false
  "filter": {
    "supported":true
  "changePassword" : {
   "supported":true
  "sort": {
    "supported":false
  "etag": {
    "supported":false
  "xmlDataFormat": {
    "supported":false
  "authenticationSchemes": [
      "name": "HTTP Basic",
     "description": "Authentication using HTTP Basic",
     "type": "httpbasic"
    },
```

```
"name": "TLS Client Certificate",
    "description": "Authentication via TLS Client Certificate",
    "type":"tls"
]
```

System-services endpoints

System-services endpoints generally apply to the PingFederate server, whether used as an identity provider (IdP), service provider (SP), or both. Parameters are case-sensitive.

/pf/heartbeat.ping

This endpoint returns an HTTP status code of 200 and a message body of **OK** if the PingFederate runtime server is up and functional. You can customize the message by modifying a PingFederate property and a Velocity template file. For more information, see Customizing the heartbeat message on page 835.



If a GET request receives a connection error or an HTTP status code other than 200, the server associated with the endpoint is down or malfunctioning.

Load balancers can use this endpoint to determine the status of PingFederate independently of checks used to determine the status of the supporting hardware.

You can also configure the server to provide regular status information to a network-management utility. For more information, see *Runtime reporting* on page 864.

/pf/adapter2adapter.ping

This endpoint initiates direct IdP-to-SP adapter mapping, when that feature is configured in the Adapterto-Adapter Mappings window. For more information, see Adapter-to-adapter mappings on page 371.



To prevent users from circumventing the SP authentication policies, this endpoint becomes inactive when SP authentication policies are enabled but IdP authentication policies are disabled. Administrators can configure SP authentication policies for the internal users to re-enable access to protected resources.

For information, see Configure SP authentication policies for internal users.

The following table shows the HTTP parameters for this endpoint.

Parameter	Description
TargetResource (optional)	Indicates where the user is redirected after a successful SSO. If this parameter is not included in the request, PingFederate redirects the user to a default
(optional)	location if one is specified in the Applications # Integration # SP Default URLs window.
InErrorResource	Indicates where the user is redirected if the SSO is unsuccessful. If this parameter is not included in the request, PingFederate redirects the user to the SSO error landing page hosted within PingFederate. For more information, see <i>Customizable user-facing pages</i> on page 798.
(optional)	

Parameter	Description
IdpAdapterId (optional)	Indicates the IdP adapter instance to use for authentication if more than one IdP adapter is configured in adapter-to-adapter mappings.
SpSessionAuthnAdapt (optional)	eladicates the SP adapter instance to be used. If not provided and more than one SP adapter instance is configured with adapter-to-adapter mapping, PingFederate selects one based on entries defined in the Applications # Integration # Target URL Mapping window. For more information, see Configuring target URL mapping on page 652.
ChangePassword	If a request includes this parameter with a value of true and invokes an HTML Form Adapter instance, the user is redirected to the Change Password template and prompted to update the network password.
	Note: To use this parameter, the Allow Password Changes check box must be selected in the adapter configuration of the invoked HTML Form Adapter instance. For more information, see .

/pf/sts.wst

This endpoint initiates direct security token service (STS) token-to-token exchange and token validation from an IdP token processor to an SP token generator, when that feature is configured in the Token **Translator Mappings** window. For more information, see *Token translator mappings* on page 377.

The following table shows the HTTP parameters for this endpoint.

Parameter	Description
TokenProcessorId	Indicates the IdP token processor to use in the mapping. Required when multiple IdP token processors are configured in token-to-token mappings.
TokenGeneratorId	Indicates the SP token generator to use in the mapping. Required when multiple SP token generators are configured in token-to-token mappings.



Important:

If mutual SSL/TLS is used for authentication, you must configure a secondary PingFederate listening port used by partners or STS clients for the relevant endpoints—*.ssaml* and *.wst. For more information, see Configuring PingFederate properties on page 752.

/pf/sts_mex.ping

This endpoint returns STS metadata for use in expediting configuration of web-service applications.

The following table shows the HTTP parameters for this endpoint.

Parameter	Description
PartnerSpId	The connection ID of the SP to whom the SAML token will be issued. This parameter determines the connection for which metadata will be generated.
PartnerIdpId	The connection ID of the IdP issuing the SAML token to be consumed by PingFederate. This parameter determines the connection for which the metadata will be generated.

Parameter	Description
vsid	Specify the virtual server ID.
(optional)	If absent, PingFederate uses the default virtual server ID (if specified) for the connection or the federation ID defined on the System # Server # Protocol Settings # Federation Info tab.



If your partner fails to retrieve metadata when sending both the PartnerSpId or the PartnerIdpId, and the vsid query parameters, perhaps it is only capable of sending one query parameter in such requests. An alternative metadata exchange endpoint that includes the virtual server ID information should resolve the issue.

For more information, see .

/pf/federation_metadata.ping

This endpoint returns SAML and WS-Federation metadata.

The following table shows the HTTP parameters for this endpoint.

Parameter	Description
PartnerSpId	The connection ID of the SP to whom the assertions or tokens are issued. This parameter determines the connection for which metadata is generated.
PartnerIdpId	The connection ID of the IdP issuing the assertions or tokens to be consumed by PingFederate. This parameter determines the connection for which the metadata is generated.
vsid	Specify the virtual server ID.
(optional)	If absent, PingFederate generates the metadata based on the connection's default virtual server ID, if two or more virtual server IDs are defined, or the federation ID defined on the System # Server # Protocol Settings # Federation Info tab.



If your partner fails to retrieve metadata when sending both the PartnerSpId or the PartnerIdpId, and the vsid query parameters, perhaps it is only capable of sending one query parameter in such requests. An alternative metadata exchange endpoint that includes the virtual server ID information should resolve the issue.

For more information, see .

Constructing an alternative metadata exchange endpoint

You can embed virtual server ID information into a security token service (STS) metadata exchange endpoint or a SAML and WS-Federation metadata exchange endpoint.

About this task

This process is useful for scenarios where partners prefer to retrieve metadata by sending one query parameter such as PartnerSpId or PartnerIdpId, instead of two query parameters such asPartnerSpId or PartnerIdpId and vsid.

Steps

 Construct a JSON object containing a key-value pair of the virtual server ID by using the following format.

```
{"vsid":"<VirtualServerIdValue>"}
```

Example:

For example, if the virtual server ID is Engineering, the JSON object is {"vsid": "Engineering"}.

2. Base64url-encode the JSON object.

Example:

For example, if the JSON object is {"vsid": "Engineering"}, the base64url-encoded value is eyJ2c21kIjoiRW5naW51ZXJpbmcifQ.

For more information about base64url, see tools.ietf.org/html/rfc4648.

3. Insert the base64url-encoded value prefixed with a forward slash into the metadata exchange endpoints, described as follows:

Federation metadata endpoint (/pf/federation metadata.ping)

```
Between /pf and /federation metadata.ping.
```

STS metadata endpoint (/pf/sts mex.ping)

Between /pf and /sts mex.ping.

Example:

For example, if the base64url-encoded value is eyJ2c2lkIjoiRW5naW5lZXJpbmcifQ, the metadata exchange endpoints embedding with the virtual server ID are:

Federation metadata endpoint

```
/pf/eyJ2c2lkIjoiRW5naW5lZXJpbmcifQ/federation metadata.ping
```

Example: https://idp.example.com:9031/pf/eyJ2c2lkljoiRW5naW5lZXJpbmcifQ/federation_metadata.ping?PartnerSpId=sp.example.org

STS metadata endpoint

```
/pf/eyJ2c2lkIjoiRW5naW5lZXJpbmcifQ/sts mex.ping
```

Example: https://idp.example.com:9031/pf/eyJ2c2lkljoiRW5naW5lZXJpbmcifQ/sts_mex.ping? PartnerSpId=sp.example.org

Authentication API

The PingFederate authentication API is a JSON-based API that enables end-user interactions, such as credential prompts, to be handled by an external web application. This API does so by providing access to the current state of the flow as an end user steps through a PingFederate authentication policy.

Authentication flows are initiated through browser-based single sign-on (SSO) application endpoints, such as /idp/startSSO.ping, or a protocol request, such as an OpenID Connect authentication request received at the authorization endpoint: /as/authorization.oauth2. As PingFederate runs the configured authentication policy, if it encounters an API-capable adapter or selector, and an authentication application is configured for the policy, PingFederate redirects to the authentication application's URL, passing the ID of the flow in the flowId query parameter.

The authentication application can then retrieve the current state of the flow by issuing a **GET** request to the $/pf-ws/authn/flows/{flowId}$ endpoint. The **_links** field in the response lists the available authentication actions that can be performed from the current state. To invoke an action, the authentication

application sends a POST request to the /pf-ws/authn/flows/{flowId} endpoint. The action ID is specified via the Content-Type HTTP request header in the format

application/vnd.pingidentity.<actionId>+json

When the application invokes an action, PingFederate responds either with the next state for the flow or an error.

When the user completes the authentication policy steps successfully, the authentication API returns a RESUME status to the authentication application. This status indicates that the API client should redirect the user's browser to the resumeUrl specified in the response. PingFederate will then be responsible for the final step in the flow, such as passing a SAML assertion to a partner. A RESUME status will also be sent if PingFederate encounters an identity provider (IdP) connection in the policy tree, or an IdP adapter or selector that is not API-capable. When the API client redirects the user, PingFederate will take the steps needed to invoke the authentication source.

If the user has interacted with an authentication application and the flow terminates with an error, the API client will receive a FAILED status from the API.



To avoid issues with third-party cookies in some browsers, you should give the authentication application the same parent domain as the PingFederate authentication API URL that the application accesses. This could be a common domain of PingFederate's base URL or one of PingFederate's defined virtual hosts.

Key concepts

Flow

The SSO transaction invoking the authentication API.

States

The available states (if any) for a given API-enabled adapter or selector.

Current state

Indicates what the adapter or selector is ready to do next.

Actions

The available actions (if any) for a given state.

Session management

The authentication API endpoint has the same domain as PingFederate's other application endpoints and shares the PingFederate session cookies with those endpoints. This ensures that session data created by API applications can be retrieved when the user interacts with PingFederate's other endpoints and vice versa.



Note:

Because the authentication API relies on the PingFederate session cookies, only browser-based web applications can currently make use of the API. Server-side web applications are not supported.

In order for PingFederate to manage session cookies correctly, the Javascript-written authentication application must set the withCredentials flag on the XMLHttpRequest object to true.

To protect against cross-site request forgeries, API clients are also required to include an X-XSRF-Header HTTP request header with each request, for example: x-xsrf-Header: PingFederate. This custom

header ensures that browsers enforce cross-origin resource sharing (CORS) policies when API requests are sent. This header can have any value.

Authentication API Explorer

PingFederate includes an API Explorer, which allows you to view the states, actions, and models available for the various API-capable adapters and selectors included in your PingFederate environment. The endpoint for the Authentication API Explorer is /pf-ws/authn/explorer.

You can download a Postman collection file from **Authentication** # **Integration** # **Authentication API Applications**. The file contains information converted from the Explorer into a Postman collection. You can then import the collection file into the Postman application. The collection file provides every possible action of every state in every plugin deployed in the PingFederate instance. It contains:

- All API-capable plugins, both adapters and selectors, with all of their states and actions
- Pre-populated body containing information about required information types and parameters
- Headers for API calls, both X-XSRF-Header and Content-Type

For more information, see *Exploring the authentication API* on page 1156.

JavaScript widget for the PingFederate authentication API

The JavaScript Widget for the PingFederate authentication API is a customizable JavaScript library that provides the capabilities of the HTML Form Adapter and Identifier First Adapter through the authentication API:

- User Login
- Trouble Signing in
- Trouble with Username
- Password Reset

The widget is a ready-to-use drop-in bundle with CSS and customizable templates. This alternative to the PingFederate templates provides a sign-in experience as a single page application. For more information, go to *JavaScript Widget for the PingFederate Authentication API*.

Exploring the authentication API

PingFederate includes an authentication API explorer that allows you to view the states, actions, and models available for the various API-capable adapters and selectors included in your PingFederate environment.

Steps

Enable Authentication API

- 1. Go to Authentication # Integration # Authentication API Applications.
- 2. Select the Enable Authentication API check box.
- 3. Select the Enable API Explorer check box.

Configure an authentication application for the Authentication API Explorer

- 4. Go to Authentication # Integration # Authentication API Applications.
- 5. Click Add Authentication Application.
- 6. On the Authentication Application window, configure each field as described in the following table.

Field	Description
Name	A name of the authentication application, such as Authentication API Explorer

Field	Description
Description	An optional description of the authentication application, such as Explore the authentication API!
URL	A combination of the PingFederate base URL and the application path / pf-ws/authn/explorer
	For example, if the base URL is https://localhost:9031, enter: https://localhost:9031/pf-ws/authn/explorer
	You can find your PingFederate base URL at System # Server # Protocol Settings # Federation Info .
Additional Allowed Origins	Any additional allowed origins. For more information, see <i>Configuring</i> authentication applications on page 393.
	If you are using a PingFederate base URL of https://localhost:9031 for testing purpose, you can skip this field.

- 7. On the Authentication Application window, click Save.
- 8. On the Authentication API Applications window, click Save.

Explore the available states, actions, and models for any API-capable adapter or selector in your PingFederate environment

9. Browse to the URL of the Authentication API Explorer, as in the following example. Example:

https://localhost:9031/pf-ws/authn/explorer

10. From the **Authentication Adapter/Selector** list, select an authentication adapter or selector.

Result:

Based on your selection, the Authentication API Explorer displays a list of states. You can then inspect the following items for any given state:

- The state purpose
- The state data model (if any)
- The available action of actions (if any)
- The action data model (if any) for a given action
- The errors (if any) for a given action

The developers of your web applications use this information to create the desired authentication experience.

Explore the authentication API through a request without an authentication policy

11. Configure a use case to use an instance of the HTML Form Adapter for authentication.

For example, you can create an SP connection that uses an instance of the HTML Form Adapter for authentication. For more information, see *Configuring a sample use case* on page 236.

- 12. Go to Authentication # Integration # Authentication API Applications.
- 13. From the **Default Authentication Application** list, select the authentication application that represents the Authentication API Explorer. Click **Save**.
- 14. Initiate a request supported by the use case you configured in step 11.

Result:

When PingFederate receives your request, it gathers from your use case that it should invoke the HTML Form Adapter. Because the HTML Form Adapter is API-enabled and you have configured the Authentication API Explorer to be the default authentication application, instead of returning the **Sign On** window (from the HTML Form Adapter), PingFederate redirects the browser to the Authentication

API Explorer with a flowid query parameter. For example: https://localhost:9031/pf-ws/authn/explorer?flowId=Tt9n7.

The Authentication API Explorer opens and pre-populates the **Flow ID** field with the flow ID value generated by PingFederate, Tt9n7 in the above example.

15. In the Authentication API Explorer, click **Get** next to the pre-populated flow ID value.

Result:

The Authentication API Explorer displays a JSON response as the result of the GET request. This response contains information that the web application requires to proceed further. For instance, the status parameter value indicates the current state of the request. Because the sample request invokes the HTML Form Adapter, the current state should be USERNAME_PASSWORD_REQUIRED.

At the end of the result is a hyperlink to the current state. In this example, when you select the current state link, the Authentication API Explorer jumps to the **USERNAME_PASSWORD_REQUIRED** state and expands its contents for further review.

From this point, you can review the state data model and move the request further by selecting the appropriate action and action data, if it is required.

Explore the authentication API through a request with an authentication policy

- 16. Go to Authentication # Integration # Authentication API Applications.
- 17. From the **Default Authentication Application** list, select **Select**.

Result:

As a result, PingFederate does not designate a default authentication application.

- 18. Define an authentication policy to use the Authentication API Explorer:
 - a. On Authentication # Policies # Policies, click Add Policy.
 - b. Enter a policy name and optionally a description.
 - c. From the **Authentication Application** list, select the authentication application that represents the Authentication API Explorer.
 - d. Select the HTML Form Adapter instance that has been mapped to your use case in step 11.

For both the **Fail** and **Success** policy paths, select **Done**.

- e. Click Done.
- f. Select the IdP Authentication Policies check box.
- g. Click Save.
- 19. Initiate a request supported by the use case you configured in step 11.
- 20. Use the Authentication API Explorer to learn more about the authentication API.

Generate a Postman collection file

21. Browse to the URL of the Authentication API Explorer as in the following example.

Example:

https://yourpingfederatebaseURL:xxxx/pf-ws/authn/explorer

- 22. Click the orange **Postman Collection** button.
- 23. Navigate to the location where you want to save the postman_collection.json file. Then click Save.
- 24. Open Postman and import the file.

- 25. In Postman, manually configure the following:
 - The flowid collection variable. When PingFederate receives a request, it generates the flowid and displays it in the Flow ID field on the Authentication API Explorer window; for example, Tt9n7.
 - The baseUrl collection variable. This is the base URL of your Authentication API Explorer; for example: https://localhost:9031.
 - The PingFederate cookie.

You must also modify the body, if one exists, to ensure that API calls work correctly.

Mobile application authentication through REST APIs

In PingFederate, you can configure mobile applications to authenticate through REST APIs as OAuth clients without needing to handle HTTP redirections. When authentication is complete, the applications receive an OAuth authorization code or access token, and possibly an OpenID Connect ID token.

Single-page web applications can also use redirectless mode if administrators configure them in PingFederate as authentication applications. For more information, see *Configuring authentication* applications on page 393. Web-based authentication applications must be highly trusted. See Denying authentication applications access to the authorization endpoint on page 394 for more information on security considerations for authentication applications.

Administrators can allow an OAuth client to initiate authorization directly through the authentication API:

- 1. Go to the client's configuration in the Client window.
- 2. Select the Allow Authentication API OAuth Initiation check box.

When enabling this feature, consider the following:

- Redirection URLs are optional, but without a redirection URL, browser-based OAuth flows will not work.
- This flow does not support the user-facing scope consent page, Request for Approval. So, enabling this feature automatically enables the Bypass Authorization Approval feature and Restrict Common Scopes feature.
- The application must manage the PF cookie and, if persistent authentication sessions are configured, the PF.PERSISTENT cookie.

To authenticate with this method, an OAuth client makes two or more API requests:

1. The OAuth client initiates authentication by calling the OAuth 2.0 authorization endpoint /as/ authorization.oauth2 to get a flow ID and other information it needs for the next request. The client must specify the pi.flow response mode in this call. See the first request sample below.



Note:

If a valid authentication session already exists when the OAuth client makes the first request, the client will receive a response with a token. In this case, the client does not need to make the next request.

2. The client calls the authentication API flow endpoint /pf-ws/authn/flows/{flow id} to get the token. See the second request sample below.



In some cases, depending on the configuration of the authentication policy, the client must make more than two requests to get a token.

First request sample

```
GET /as/authorization.oauth2?
client_id=im_client&response_type=token&response_mode=pi.flow HTTP/1.1
Host: www.example.com:9031
X-XSRF-Header: PingFederate
```

Sample response 1 to the first request: If a valid authentication session does not already exist, the OAuth client receives a response that provides the information the client needs for the next request (see the second request sample). The <code>self</code> link in the response shows that the second request must call the <code>pf-ws/authn/flows</code> endpoint, rather than the authorization endpoint. The value of the <code>id</code> field is the ID of the flow that was created. The client must append the flow ID to the <code>pf-ws/authn/flows</code> endpoint for subsequent API requests.

Sample response 2 to the first request: If a valid authentication session already exists, the OAuth client receives a response with a token. So the client does not need to make the second request.

```
"id": "PyH5g", // Flow ID
   "pluginTypeId": "7RmQNDWaOnBoudTufx2sEw",
   "status": "COMPLETED",
   "authorizeResponse": {
        "access_token": "000144Qlv9eqpBk03ngAd7M35Gaj41Mgisk",
        "token_type": "Bearer"
   "_links": {
        "self": {
            "href": "https://www.example.com:9031/pf-ws/authn/flows/PyH5g"
        }
}
```

Second request sample

```
POST /pf-ws/authn/flows/PyH5g HTTP/1.1
Host: www.example.com:9031
X-XSRF-Header: PingFederate
Content-Type: application/vnd.pingidentity.checkUsernamePassword+json
Cookie: PF=8PgutwFizNS7EoaiB0qVsa

{
   "username": "joe",
```

```
"password": "Password1"
}
```

Sample response to the second request

Development of authentication API-capable adapters and selectors

PingFederate's authentication API lets applications interact with authentication policies. Making an adapter or selector plugin API-capable means ensuring that an authentication application can invoke the plugin through this API.

API-capable plugins must handle JSON-based API requests. When a plugin is invoked through the authentication API endpoint, if it needs interaction from the user, the plugin sends a JSON-based response rather than rendering a template.

Adapter and selector plugins handle distinct kinds of requests and generate distinct kinds of responses. The main method you implement in adapters is lookupAuthN(). The main method you implement in selectors is selectContext().

Developing an API-capable plugin requires a dependency on the PingFederate authentication API SDK JAR file, pf-authn-api-sdk-version.jar. In the PingFederate installation package, you can find the SDK JAR file in the server/default/lib directory. Documentation for the classes are in the Javadocs for the standard PingFederate SDK, under sdk/doc/index.html in the PingFederate install package.

Authentication API states, actions, and models

To develop authentication API-capable adapters and selectors, you must understand the states, actions, and models of single sign-on (SSO) transactions through the PingFederate authentication API.

PingFederate assigns a flow ID to each SSO transaction that uses the authentication API. PingFederate uses the flow ID to determine a transaction's state.

As a user steps through an SSO transaction, the transaction is always in some state. The state includes a status field and other fields specific to that state. The class containing those other fields is the model for the state.

The API endpoint returns the following when the user's SSO transaction has reached the USERNAME PASSWORD REQUIRED state for the form adapter.

```
"id": "PyH5g",
"pluginTypeId": "7RmQNDWaOnBoudTufx2sEw",
"status": "USERNAME_PASSWORD_REQUIRED",
"showRememberMyUsername": false,
"showThisIsMyDevice": false,
"thisIsMyDeviceSelected": false,
```

The model for this state is the class UsernamePasswordRequired. It includes fields such as showThisIsMyDevice, which help the API client know how to render the credential prompt to the user.

The API response also includes a list of available actions. In this case, the only action available is checkUsernamePassword. The API client can select this action by sending a POST request with the Content-Type of application/vnd.pingidentity.checkUsernamePassword +json. Each action has its own model containing the fields that the POST body can provide. For the checkUsernamePassword action, the model is CheckUsernamePassword.

The POST body can be as simple as the following.

```
"username": "joe",
 "password": "2Federate"
}
```

After receiving this request, PingFederate calls the lookupAuthN() method of the form adapter. If the form adapter encounters an error while validating the credentials, it writes a JSON API error to the response. If the form adapter successfully validates the credentials, it returns AUTHN STATUS.SUCCESS from its lookupAuthN() method. PingFederate then goes to the next step in the authentication policy. If the next step is an API-capable adapter, PingFederate calls 100kupAuthN () on that adapter and the adapter determines its current state and writes it to the response, along with the available actions.



Note:

The PingFederate authentication API follows a different naming convention for actions than PingOne. PingOne names actions as noun. verb, such as otp.check. PingFederate uses verbNoun, such as checkOtp.

Specification of the plugin API

The first step in adding API support to your plugin is to implement the AuthnApiPlugin interface.

The AuthnApiPlugin interface has two methods: qetApiSpec() and qetApiPluginDescriptor(). You only need to implement the getApiSpec() method. The API specification this method returns defines the states, models, and actions that your plugin exposes in the API.

The API specification is defined by the *Spec classes in the SDK. These include AuthnStateSpec, AuthnActionSpec, AuthnErrorSpec, and AuthnErrorDetailSpec. The information in these classes lets the PingFederate authentication API Explorer provide documentation for API client developers. That documentation describes your plugin's API and lets developers experiment with it.

You can access the API Explorer at https://PingFederate_host:9031/pf-ws/authn/explorer. To enable the API Explorer, go to the Authentication API Applications window and select the Enable API Explorer check box. An easy way to use the API Explorer is to create an authentication API application in PingFederate and set the URL for the application to the API Explorer's URL.

Note:

When defining models for states and actions, use the @Schema annotation to describe each field in the model and show whether the field is required.

The rest of this document primarily uses the TemplateRenderAdapter as an example. The source for this adapter is in the PingFederate installation package's sdk/plugin-src/template-renderadapter-example directory. This adapter is simple. It just prompts the user to enter their username and provide a set of string attributes. The administrator defines the list of attributes by extending the adapter contract. The attribute values are passed back in the SubmitUserAttributes model as a map. Representing field values using a map in the model is unusual. Usually a separate field in the model defines each allowed field, which provides better type safety in the code.

State model example

The following is an example of a state model used by the TemplateRenderAdapter.

```
/**
* This is the model for the USER ATTRIBUTES REQUIRED state. It defines the
* fields that are returned to the API client in a GET response for this
state.
public class UserAttributesRequired
private List<String> attributeNames = new ArrayList<>();
 * Get the list of user attributes supported by this adapter instance.
  * It is recommended to annotate each getter with the @Schema annotation
  * and provide a description. This description will be used in
   generating API documentation.
  */
 @Schema(description="A list of user attribute names that are supported by
 this adapter instance.")
public List<String> getAttributeNames()
 return attributeNames;
 }
 /**
 * Set the list of user attributes supported by this adapter instance.
public void setAttributeNames(List<String> attributeNames)
 this.attributeNames = attributeNames;
 }
```

Action model example

The following is the model for the submitUserAttributes action.

```
/**
* This is the model for the submitUserAttributes API action. It defines
 * fields that may be included in the POST body for this action.
public class SubmitUserAttributes
private String username;
private Map<String,Object> userAttributes = new HashMap<>();
```

```
/**
* Get the username.
 * It is recommended to annotate each getter with the @Schema annotation
* and provide a description. The 'required' flag can also be specified.
This
 * information will be used in generating API documentation.
 * /
@Schema(description="The user's username.", required=true)
public String getUsername()
return username;
}
/**
* Set the username.
public void setUsername(String username)
 this.username = username;
* Get the user attributes.
@Schema(description="Additional user attributes, as name-value pairs.")
public Map<String, Object> getUserAttributes()
 return userAttributes;
* Set the user attributes.
public void setUserAttributes(Map<String, Object> userAttributes)
 this.userAttributes = userAttributes;
```

Action specification example

The following is the specification for the submitUserAttributes action.

```
public final static AuthnActionSpec<SubmitUserAttributes>
SUBMIT_USER_ATTRIBUTES = new
AuthnActionSpec.Builder<SubmitUserAttributes>()
.id("submitUserAttributes")
.description("Submit the user's username and attributes.")
.modelClass(SubmitUserAttributes.class)
.error(CommonErrorSpec.VALIDATION_ERROR)
.errorDetail(ErrorDetailSpec.INVALID_ATTRIBUTE_NAME)
.build();
```

Error specifications

Action specifications can include a list of possible errors and error details. Each top-level error that an authentication API request returns can include one or more error detail objects underneath it.

Typically, in the API specification, your only top-level error will be <code>CommonErrorSpec.VALIDATION_ERROR</code>. However, you can include error detail specifications that can appear under that top-level error.

```
public final static AuthnErrorDetailSpec INVALID ATTRIBUTE NAME = new
AuthnErrorDetailSpec.Builder()
 .code("INVALID ATTRIBUTE NAME")
 .message("An invalid attribute name was provided.")
 .parentCode(CommonErrorSpec.VALIDATION ERROR.getCode())
 .build();
```

Note:

The error detail specification must reference the error code of its parent top-level error. This ensures that the authentication API Explorer correctly represents the error information.

INVALID ATTRIBUTE NAME is an example of an error that would be useful for API client developers but not for end users.

For more information about defining user-facing errors, see Error messages and localization on page 1169.

State model contents

The model for a state includes all the information an API client would need to build a form (not necessarily an HTML form) to show the user. A state model should not include the text for messages to display to the user.

Defining messages for users, and localizing them if needed, is the responsibility of the API client. One of the reasons we avoid including messages in state models is that those messages will often end up including semantic content that the API client needs to drive its code. Following the rule that models do not include messages helps ensure that our models include all the fields that an API client needs to provide the desired user experience.

Sometimes following this rule requires you to add more states. This is preferable to embedding the state information inside of a message because it makes it easier for an API client to handle that state in the desired way.

The one exception we have to this rule is around error messages. API errors include error message text, and in some cases, API clients will display the message text directly to users. This avoids every API client having to write its own messages for every user-facing error the API can generate. For more information, see Error messages and localization on page 1169.

Non-interactive plugins

Some plugins, typically selectors, do not need to interact with the user to do their job. Making these plugins API-capable is straight-forward.

You still implement the AuthnApiPlugin interface, but you can just return null from the getApiSpec() method. And then you override the default implementation of getApiPluginDescriptor() and return an AuthnApiPluginDescriptor instance with the interactive flag set to false. As with many other classes in the SDK, there is an AuthnApiPluginDescriptor.Builder class to help in creating the descriptor.

When interactive is false, PingFederate knows that it never needs to redirect when it encounters your selector. If the request is occurring on the API endpoint, PingFederate can immediately call selectContext(). The same is true if the request is occurring on a front-channel endpoint, such as / as/authorization.oauth2.

If your selector does not implement AuthnApiPlugin, then PingFederate assumes that only a frontchannel endpoint can call your selector. If PingFederate encounters your selector while executing an

API request, PingFederate will send a RESUME response to the API client so that the user is redirected to PingFederate.

Runtime behavior implementation

After you specify your plugin's API at least partially, you can start implementing the runtime behavior. Use the specification that you defined previously to implement the runtime functionality.

Follow this pattern in lookupAuthN():

- 1. Check for the possible actions the adapter expects in the current state.
- 2. If an action is matched, then try to extract the expected model from the request and handle the action.
- 3. If an action is requested, but it does not match an action allowed for the current state, then return an INVALID ACTION ID error.
- 4. If no action is requested, render the response for the current state.

The AuthnApiSupport class provides much of the functionality for handling API requests and sending responses. The TemplateRenderAdapter stores a reference to this singleton in its apiSupport field.

```
private AuthnApiSupport apiSupport = AuthnApiSupport.getDefault();
```

Checking for actions

The following code shows the preferred approach for checking for the submitIdentifiers action.

The adapter performs this check two ways, depending on whether the current request is from the API endpoint. The TemplateRenderAdapter uses a slightly different but equivalent method.

```
/**
  * Determine if the user chose "Submit".
  */
private boolean isSubmitAttributesRequest(HttpServletRequest req)
{
  if (apiSupport.isApiRequest(req))
  {
    return ActionSpec.SUBMIT_USER_ATTRIBUTES.isRequested(req);
  }
  return StringUtils.isNotBlank(req.getParameter("pf.submit"));
}
```

Extracting models from requests

The next step extracts the model from the request. This step varies depending on whether the request is from the API endpoint. For an API request, call the <code>AuthnApiSupport.deserializeAsModel()</code> method. For a non-API request, you must build the model from the parameters in the request.

```
private SubmitUserAttributes getSubmittedAttributes(HttpServletRequest req)
  throws AuthnErrorException, AuthnAdapterException
{
   if (apiSupport.isApiRequest(req))
   {
     try
     {
       return apiSupport.deserializeAsModel(req, SubmitUserAttributes.class);
   }
   catch (IOException e)
   {
      throw new AuthnAdapterException(e);
   }
} else
{
   SubmitUserAttributes result = new SubmitUserAttributes();
```

The deserializeAsModel () method also does some validation on the incoming JSON. This includes checking for fields flagged as required in the model, using the @Schema annotation. If a validation error occurs during this step, the method throws an AuthnErrorException, which the adapter can convert to an API error response. For more information, see *Handling authentication error* exceptions on page 1168.

Performing additional validation

The deserializeAsModel() method performs some basic validation on the submitted JSON. Your adapter probably needs to perform more validation and send an AuthnError to the API client if it finds any errors. Here is how the TemplateRenderAdapter validates the names of the provided user attributes:

Handling invalid action IDs

If a request from an API client includes an action ID that does not match any actions available in the current state, it is best practice to return an error to the client.

After checking all the possible actions, if none match and the request's action ID is not null, the adapter can throw an AuthnErrorException. The adapter catches this exception and writes an error to the API response.

```
if (apiSupport.getActionId(req) != null)
{
   // An action ID was provided but it does not match one of those expected in
   the current state.
   throw new
   AuthnErrorException(CommonErrorSpec.INVALID_ACTION_ID.makeInstance());
}
```

If the <code>deserializeAsModel()</code> method detects an error while deserializing the model, it throws an <code>AuthnErrorException</code>. If the added validation checks in <code>validateSubmittedAttributes</code> detect an error, they also throw this exception.

The adapter should catch this exception and send an API error response using the method AuthnApiSupport.writeErrorResponse().

```
try
{
...
}
catch (AuthnErrorException e)
{
   // A validation error occurred while processing an API request, return an error response to the API client
   apiSupport.writeErrorResponse(req, resp, e.getValidationError());
   authnAdapterResponse.setAuthnStatus(AUTHN_STATUS.IN_PROGRESS);
   return authnAdapterResponse;
}
```

Sending API responses

AuthnApiSupport provides several methods for writing API responses.

The following example shows how the TemplateRenderAdapter writes the response for the USER ATTRIBUTES REQUIRED state.

```
private void renderApiResponse(HttpServletRequest req, HttpServletResponse
  resp, Map<String, Object> inParameters) throws AuthnAdapterException
{
   UserAttributesRequired model = new UserAttributesRequired();
   model.setAttributeNames(new ArrayList<>(extendedAttr));
   AuthnState<UserAttributesRequired> authnState =
   apiSupport.makeAuthnState(req, StateSpec.USER_ATTRIBUTES_REQUIRED, model);
   try
   {
     apiSupport.writeAuthnStateResponse(req, resp, authnState);
   }
   catch (IOException e)
   {
     throw new AuthnAdapterException(e);
   }
}
```

The <code>makeAuthnState()</code> method takes an <code>AuthnStateSpec</code> and an instance of the model for that state and builds an <code>AuthnState</code> object. The <code>AuthnState</code> object can then be further modified. For example, you could remove an action that is not currently applicable using the <code>removeAction()</code> method. Then you write the <code>AuthnState</code> object to the response using the <code>writeAuthnStateResponse()</code> method.

Returning authentication statuses

As with non-API requests, when the adapter finishes, it returns <code>AUTHN_STATUS.SUCCESS</code> or <code>AUTHN_STATUS.FAILURE</code> from <code>lookupAuthN()</code>.

If the adapter has not yet finished and has written something to the response, it should return AUTHN_STATUS.IN_PROGRESS.

Session state management

Session state management for authentication API adapters is no different than for regular adapters. The same mechanisms, such as <code>SessionStateSupport</code> and <code>TransactionalStateSupport</code>, are used to store and retrieve session attributes on behalf of a user.

It should not be necessary to store more state attributes just to support the authentication API. The same session attributes should cover both API and non-API requests.

You can also wrap an API-capable adapter in a PingFederate-managed authentication session. PingFederate-managed authentication sessions mean that many adapters no longer need to provide their own internal session tracking.

Error messages and localization

Error messages are the one case where an API response could include user-facing text. The typical case is a validation error.

For validation errors, the adapter constructs an AuthnError with the code VALIDATION_ERROR, and then adds AuthnErrorDetail objects for each of the errors that occurred. The userMessage field of the AuthnErrorDetail object provides the user-facing text. Like states and actions, you can define errors up front using an AuthnErrorSpec or an AuthnErrorDetailSpec. Then an instance of the error is constructed from the specification on demand.

The following example shows how you can define the specification for an invalid OTP error.

```
public final static AuthnErrorDetailSpec INVALID_OTP = new
AuthnErrorDetailSpec.Builder()
.code("INVALID_OTP")
.message("An invalid or expired OTP code was provided.")
.userMessage("This code is invalid or has expired.")
.parentCode(CommonErrorSpec.VALIDATION_ERROR.getCode())
.build();
```

The following example shows how you can use that specification to send an error response to the API client.

```
AuthnErrorDetail errorDetail =
   ErrorDetailSpec.INVALID_OTP.makeInstanceBuilder().build();
AuthnError authnError =
   CommonErrorSpec.VALIDATION_ERROR.makeInstanceBuilder().detail(errorDetail).build();
apiSupport.writeErrorResponse(req, resp, authnError);
```

To localize the error message using a properties file for your adapter, you can use LocaleUtil and LanguagePackMessages from the standard PingFederate SDK.

```
LanguagePackMessages messages = new LanguagePackMessages("my-adapter-
messages", LocaleUtil.getUserLocale(req));
String errorMessage = messages.getMessage("invalid.otp.key", new String[]
{});
AuthnErrorDetail errorDetail =
   ErrorDetailSpec.INVALID_OTP.makeInstanceBuilder().userMessage(errorMessage).build();
AuthnError authnError =
   CommonErrorSpec.VALIDATION_ERROR.makeInstanceBuilder().detail(errorDetail).build();
apiSupport.writeErrorResponse(req, resp, authnError);
```

For more information about how PingFederate determines the user's locale at runtime, see *Locale overrides by cookies* on page 818.

Some errors reflect problems with API client programming rather than with end user input. If you think an error will not be shown to an end user, then you do not need to populate the userMessage field.

Legal Information

PingFederate Server documentation

© 2020 Ping Identity Corporation. All rights reserved.

Trademarks

Ping Identity, the Ping Identity logo, PingAccess, PingFederate, PingID, and PingOne are registered trademarks of Ping Identity Corporation ("Ping Identity"). All other trademarks or registered trademarks are the property of their respective owners.

Disclaimer

The information provided in these documents is provided "as is" without warranty of any kind. Ping Identity disclaims all warranties, either express or implied, including the warranties of merchantability and fitness for a particular purpose. In no event shall Ping Identity or its suppliers be liable for any damages whatsoever including direct, indirect, incidental, consequential, loss of business profits or special damages, even if Ping Identity or its suppliers have been advised of the possibility of such damages. Some states do not allow the exclusion or limitation of liability for consequential or incidental damages so the foregoing limitation may not apply.

Index

U Α Assertions 29 uninstalling 126 unique values, for pseudonym creation 385 attribute sources custom 879 C web services certificates SSO Directory 1090 revocation of 623 CIAM 99 Customer IAM 99 D datastore LDAP configuration 871 directory service SOAP example 1092 Ε Exchange admin center 367 installation federation server 120 L LDAP configuration 871 SSL 871 using SSL 871 0 Office 365 Outlook web app 367 OpenID Connect 58 Ρ pseudonyms unique values for 385 S SSL using for LDAP 871 SSO 1090 SSO directory service 1090