

PingFederate[®]

Server 8.4.4



Contents

Copyright.....	8
Get started with PingFederate Server.....	9
Introduction to PingFederate.....	9
About identity federation and SSO.....	9
Security token service.....	10
OAuth authorization server.....	10
User account management.....	11
Enterprise deployment architecture.....	11
Additional features.....	11
Installation.....	12
Deployment options.....	13
System requirements.....	14
Port requirements.....	17
Install Java.....	19
Installation options.....	19
Open the PingFederate administrative console.....	21
Initial setup wizard.....	22
Install PingFederate as a service.....	25
Uninstall PingFederate.....	27
The administrative console.....	28
Tasks and steps.....	28
Console buttons.....	29
Supported standards.....	29
Federation roles.....	30
Terminology.....	30
Browser-based SSO.....	32
Web Services standards.....	46
OAuth 2.0.....	48
System for Cross-domain Identity Management (SCIM).....	50
Transport and message security.....	51
Supported hardware security modules.....	51
Install and configure Luna SA client and PingFederate.....	52
Install and configure nShield Connect client and PingFederate.....	54
Administrator's manual.....	57
Key concepts.....	57
Connection types.....	58
About WS-Trust STS.....	58
About OAuth.....	61
SSO integration kits and adapters.....	67
Hierarchical plug-in configurations.....	69
Identity mapping.....	69
About attributes.....	70
Security infrastructure.....	75
Using Auto-Connect.....	79
User provisioning.....	81

Federation planning checklist.....	82
Federation hub use cases.....	86
System administration.....	90
Manage PingFederate license.....	91
Configure PingFederate properties.....	93
PingFederate log files.....	95
Account management.....	108
Alternative console authentication.....	112
Manage email configuration.....	114
Configure virtual host names.....	115
Export metadata to an XML file.....	116
Sign XML files.....	118
Replicate configuration.....	118
Configuration archive.....	118
Automating configuration migration.....	120
Outbound provisioning CLI.....	124
Customizable user-facing screens.....	126
Customizable email notifications.....	133
Customizable text message.....	138
Localize messages for end users.....	139
Configure password policy.....	140
Manage expired OAuth persistent grants.....	140
Specify the domain of the PF cookie.....	141
Extend the lifetime of the PF cookie.....	142
Configure forward proxy server settings.....	142
Add custom HTTP response headers.....	143
Configure validation for the AudienceRestriction element.....	143
Customize the heartbeat message.....	144
Customize the favicon for application and protocol endpoints.....	144
Configure the behavior of searching multiple data stores with one mapping.....	144
System settings.....	145
Server settings.....	145
Connect to PingOne from Server Configuration.....	157
Manage data stores.....	158
IdP discovery.....	174
Configure redirect validation.....	176
Security management.....	178
Certificate management.....	178
Authentication.....	195
Self-service password reset and account unlock.....	206
Account lockout protection.....	208
Authentication policies.....	209
Selectors.....	209
Policies.....	216
Policy contracts.....	237
Adapter Mappings.....	238
Sessions.....	240
OAuth configuration.....	241
Configure OAuth use cases.....	242
Enable the OAuth AS role.....	242
Configure authorization server settings.....	243
Define scopes.....	244
OAuth access token management.....	245
Manage IdP adapter mappings for OAuth.....	253
Manage authentication policy contract mappings.....	256
Manage resource owner credential mappings.....	257

Manage access token mappings.....	260
Configure an OAuth assertion grant IdP connection.....	262
Configure OpenID Connect policies.....	266
Manage OAuth clients.....	269
Client session management.....	273
OAuth attribute mapping using a data store.....	274
Identity provider SSO configuration.....	275
IdP application integration settings.....	275
View IdP protocol endpoints.....	281
Manage SP connections.....	282
Define SP affiliations.....	327
Configure SP Auto-Connect.....	329
Service provider SSO configuration.....	330
SP application integration settings.....	331
Federation settings.....	336
Manage IdP connections.....	339
OpenID Connect Relying Party support.....	390
Configure IdP Auto-Connect.....	398
WS-Trust STS configuration.....	400
Server settings.....	400
IdP configuration for STS.....	401
SP configuration for STS.....	412
IdP-to-SP bridging.....	419
Adapter-to-adapter mappings.....	420
Token translator mappings.....	423
Bundled adapters.....	427
Kerberos Adapter.....	427
HTML Form Adapter.....	430
HTTP Basic Adapter.....	438
OpenToken Adapter.....	440
Composite Adapter.....	445
Application endpoints.....	448
IdP endpoints.....	448
SP endpoints.....	451
System-services endpoints.....	460
OAuth 2.0 endpoints.....	463
Authorization endpoint.....	464
Token endpoint.....	467
Introspection endpoint.....	472
Token revocation endpoint.....	474
Grant-management endpoint.....	475
OpenID Provider configuration endpoint.....	476
Web Service interfaces.....	477
Connection Management Service.....	477
SSO Directory Service.....	480
SOAP request and response examples.....	482
OAuth Client Management Service.....	483
OAuth Access Grant Management Service.....	488
Session Revocation API endpoint.....	490
PingFederate administrative API.....	492
Attribute mapping expressions.....	497
Enable and disable expressions.....	497
Construct OGNL expressions.....	498
Using the OGNL edit screen.....	501
Customize assertions and authentication requests.....	501
Message types and available variables.....	501

Sample customizations.....	504
Fulfillment by data store queries.....	506
Attribute mapping with multiple data sources.....	506
Data store query configuration.....	508
Troubleshooting.....	513
Enable debug messages and console logging.....	513
Server startup.....	514
Data store issues.....	514
Resolve URL-related errors.....	515
Resolve service-related errors.....	515
Troubleshoot runtime errors.....	516
Troubleshooting OAuth transactions.....	518
Other runtime issues.....	522
Glossary.....	522
List of acronyms.....	531
Server clustering guide.....	533
Overview of clustering.....	533
Deploy cluster servers.....	534
Cluster protocol architecture.....	534
Runtime state-management architectures.....	539
Runtime state-management services.....	544
Deploy provisioning failover.....	549
Configuration synchronization.....	550
Console configuration push.....	551
Configuration-archive deployment.....	551
SSO integration overview.....	552
Integration introduction.....	552
SSO integration concepts.....	552
Identity provider integration.....	553
Service provider integration.....	555
Summary.....	557
SDK developer's guide.....	559
Preface.....	559
SDK introduction.....	559
Get started with the SDK.....	561
Directory structure.....	561
Set up your project.....	561
Implementation guidelines.....	562
Shared interfaces.....	562
Implement an IdP adapter.....	563
Implement an SP adapter.....	565
Implement a token processor.....	567
Implement a token generator.....	567
Implement an authentication selector.....	567
Implement a custom data source.....	568
Implement a password credential validator.....	569
Implement an identity store provisioner.....	570
Build and deploy your project.....	575

Release notes.....	578
PingFederate 8.4.4 - November 2017.....	578
Upgrade considerations.....	578
Upgrade considerations introduced in PingFederate 7.x.....	580
Upgrade considerations introduced in PingFederate 6.x.....	581
Deprecated features.....	582
Known issues and limitations.....	583
Previous releases.....	585
PingFederate 8.4.3 - October 2017.....	585
PingFederate 8.4.2 - September 2017.....	586
PingFederate 8.4.1 - August 2017.....	587
PingFederate 8.4 - June 2017.....	588
PingFederate 8.3.3 - May 2017.....	591
PingFederate 8.3.2 - March 2017.....	592
PingFederate 8.3.1 - January 2017.....	592
PingFederate 8.3 - December 2016.....	593
PingFederate 8.2.2 - October 2016.....	596
PingFederate 8.2.1 - September 2016.....	597
PingFederate 8.2 - August 2016.....	598
PingFederate 8.1.4 - July 2016.....	601
PingFederate 8.1.3 - May 2016.....	602
PingFederate 8.1.2 - April 2016.....	603
PingFederate 8.1.1 - February 2016.....	604
PingFederate 8.1 - January 2016.....	604
PingFederate 8.0.4 - November 2015.....	607
PingFederate 8.0.3 - September 2015.....	608
PingFederate 8.0.2 - August 2015.....	609
PingFederate 8.0.1 - July 2015.....	609
PingFederate 8.0 - June 2015.....	609
PingFederate 7.3 - January 2015.....	611
PingFederate 7.2 R2 - September 2014.....	612
PingFederate 7.2.1 - August 2014.....	612
PingFederate 7.2 - June 2014.....	612
PingFederate 7.1.4 - June 2014.....	613
PingFederate 7.1 R3 - March 2014.....	613
PingFederate 7.1.3 - February 2014.....	614
PingFederate 7.1.2 - December 2013.....	614
PingFederate 7.1 R2 - December 2013.....	614
PingFederate 7.1.1 - November 2013.....	614
PingFederate 7.1 - August 2013.....	614
PingFederate 7.0.1 - May 2013.....	615
PingFederate 7.0 - April 2013.....	615
PingFederate 6.11 - December 2012.....	615
PingFederate 6.10.1 - January 2013.....	616
PingFederate 6.10 - September 2012.....	616
PingFederate 6.9 - June 2012.....	616
PingFederate 6.8 - April 2012.....	616
PingFederate 6.7 - February 2012.....	616
PingFederate 6.6 - December 2011.....	616
PingFederate 6.5.2 - November 2011.....	617
PingFederate 6.5.1 - October 2011.....	617
PingFederate 6.5 - August 2011.....	617
PingFederate 6.5-Preview - April 2011.....	617
PingFederate 6.4.1 - February 2011.....	617

PingFederate 6.4 - December 2010.....	617
PingFederate 6.3 - August 2010.....	618
PingFederate 6.3-Preview - April 2010.....	618
PingFederate 6.2 - February 2010.....	618
PingFederate 6.1 - September 2009.....	618
PingFederate 6.1-Preview - June 2009.....	619
PingFederate 6.0 - March 2009.....	619
PingFederate 5.3 - December 2008.....	619
PingFederate 5.2 - August 2008.....	620
PingFederate 5.1.1 - July 2008.....	620
PingFederate 5.1 - April 2008.....	620
PingFederate 5.0.2 - March 2008.....	621
PingFederate 5.0.1 - January 2008.....	622
PingFederate 4.4.2 - October 2007.....	623
PingFederate 4.4.1 - June 2007.....	623
PingFederate 4.4 - May 2007.....	623
PingFederate 4.3 - March 2007.....	623
PingFederate 4.2 - December 2006.....	623
PingFederate 4.1 - October 2006.....	624
PingFederate 4.0 - June 2006.....	624
PingFederate 3.0.2 - February 2006.....	624
PingFederate 3.0.1 - December 2005.....	625
PingFederate 3.0 - November 2005.....	625
PingFederate 2.1 - July 2005.....	625
PingFederate 2.0 - February 2005.....	625
PingFederate Upgrade Utility user guide.....	626
Upgrading PingFederate.....	626
Prepare for the upgrade.....	627
Run the PingFederate Upgrade Utility.....	627
Review post-upgrade tasks.....	630
Known limitations.....	636
PingFederate performance tuning.....	637
Logging.....	637
Operating system tuning.....	637
Linux tuning.....	637
Windows tuning.....	639
Concurrency.....	639
Acceptor thread pool.....	639
Server thread pool.....	640
Connection pools to data stores.....	641
Memory.....	642
JVM heap.....	642
Garbage collectors.....	642
Young generation bias.....	643
Configure JVM options.....	644
Hardware security modules.....	649
Configuration at scale.....	649
References.....	649

Copyright

PingFederate® Server documentation

© 2017 Ping Identity® Corporation. All rights reserved.

PingFederate 8.4.4
November, 2017

Ping Identity Corporation
1001 17th Street, Suite 100
Denver, CO 80202
U.S.A.

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.

Document lifetime

Ping Identity may occasionally update online documentation between releases of the related software. Consequently, if this PDF was not downloaded recently, it may not contain the most up-to-date information. Please refer to the online documentation at docs.pingidentity.com for the most current information.

From the web site, you may also download and refresh this PDF if it has been updated, as indicated by a change on this date: **November 30, 2017**.

Get started with PingFederate Server 8.4.4

This guide provides information about getting started with Ping Identity®'s PingFederate® to deploy a secure Internet-identity platform, including single sign-on (SSO) based on the latest security and e-business standards.

This guide consists of:

- [Introduction to PingFederate](#) on page 9 — A high-level view of federated identity, secure web SSO, and PingFederate features.
- [Installation](#) on page 12 — How to install PingFederate and run the administrative console for the first time.
- [The administrative console](#) on page 28 — A primer on using the administrative console and configuration screens.
- [Supported standards](#) on page 29 — An overview of industry standards that PingFederate supports, including the Security Assertion Markup Language (SAML) and WS-Federation.
- [Supported hardware security modules](#) on page 51 — How to install and configure PingFederate with a supported HSM as part of compliance with the Federal Information Processing Standard (FIPS) 140-2.

Introduction to PingFederate

Welcome to PingFederate, Ping Identity's enterprise identity bridge. PingFederate enables outbound and inbound solutions for single sign-on (SSO), federated identity management, mobile identity security, API security, and social identity integration. Browser-based SSO extends employee, customer and partner identities across domains without passwords, using only standard identity protocols (Security Assertion Markup Language—SAML, WS-Federation, WS-Trust, and OAuth).

About identity federation and SSO

Federated identity management (or “identity federation”) enables enterprises to exchange identity information securely across domains, providing browser-based SSO. Federation is also used to integrate 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 29.

Service providers and identity providers


Identity federation standards identify two operational roles in an SSO transaction: the *identity provider* (IdP) and the *service provider* (SP). An IdP, for example, might be an enterprise that manages accounts for a large number of users who may 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.

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.

PingFederate provides complete support for both roles. Note that business processes of a single organization might encompass both SP and IdP use cases; this scenario can be handled by a single instance of PingFederate.

Federation hub

To most organizations, identity federation means negotiating and managing federation settings with partners. As the number of partners grows, so does the administrative overhead. In addition, different federation protocols may also hinder application development and SSO implementation. To remove these obstacles, PingFederate® can be configured as a Federation Hub to extend federated access across partners supporting different federation standards, SAML and WS-Federation for example, as well as to provide a centralized console to simplify SSO administration. By bridging the identity providers and service providers through the federation hub, administrators also have the option to multiplex a single connection for multiple partners, adding additional use cases and reducing administration and implementation costs.



Security token service

The PingFederate WS-Trust Security Token Service (STS) allows organizations to extend SSO identity management to Web Services. (For information about WS-Trust and the role of an STS, see [Web Services standards](#) on page 46).

The STS shares the core functionality of PingFederate, including console administration, identity and attribute mapping, and certificate security management. With PingFederate, Web Services can securely identify the end user who has initiated a transaction across domains, providing enhanced service while simultaneously ensuring appropriate information access and regulatory accountability.

PingFederate can be used 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 58 in the “Key Concepts” chapter of the PingFederate *Administrator's manual*).

OAuth authorization server

PingFederate can act as an OAuth Authorization Server (OAuth AS), allowing a resource owner to grant authorization to a client requesting access to resources protected by a Resource Server. 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 Representational State Transfer (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 SaaS provider.

(For information about OAuth and the role of an AS, see [OAuth 2.0 and PingFederate AS](#) on page 48).

The PingFederate OAuth AS can be configured independently or in conjunction with STS and browser-based SSO for either an IdP or an SP deployment. For more information, see [About OAuth](#) on page 61 in the “Key Concepts” chapter of the PingFederate *Administrator's manual*.



Note: OAuth AS capabilities may require additional licenses. For more information, please contact sales@pingidentity.com.

User account management

In an identity federation, accounts are maintained for users at the IdP site. However, an SP will often have its own set of user accounts, some of which may correspond to IdP users. The SP may also need to establish and maintain parallel accounts for remote SSO users to enforce authorization policy, customize user experience, comply with regulations, or a combination of such purposes.

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

- At an IdP site, an administrator can automatically provision and maintain user accounts for partner SPs who have implemented the System for Cross-domain Identity Management (SCIM) or, when optional plug-in SaaS Connectors are used, for selected hosted-software providers.
- At an SP site, an administrator can provision accounts within the organization automatically from SCIM-enabled IdPs or use information from SAML assertions received during SSO events.

For more information, see [User provisioning](#) on page 81 in the “Key Concepts” chapter of the PingFederate *Administrator's manual*.

Enterprise deployment architecture

With PingFederate's enterprise-deployment architecture, all protocol definitions, public key infrastructure (PKI) keys, policies, profiles, etc., are managed in a single location, eliminating the need to maintain redundant copies of these configurations and trust relationships. Furthermore, when new protocols, profiles, or use cases need to be added, you only have to configure them once to make them available to your entire organization.

PingFederate also improves security by creating a single “doorway” in your perimeter through which all identity information must travel. Using PingFederate, all of your internal users who sign on to external applications exit through this doorway, while all external users who sign on to your internal systems enter through the same doorway.

The single-doorway approach also provides 100 percent visibility to all federation activities. The extensive auditing and logging capabilities of PingFederate enable you to satisfy all of your logging-related compliance and service-level requirements from a single location, as opposed to having to acquire and consolidate disparate logs from throughout your organization.

Use Case configuration

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



Tip: For IdPs, connection templates that automatically configure many steps in the administrative console are available for several use cases, including setting up SSO connections to selected SaaS vendors. (For more information, see [Outbound provisioning for IdPs](#) on page 81 in the “Key Concepts” chapter of the PingFederate *Administrator's manual*).

Additional features

PingFederate's lightweight, stand-alone architecture means you can receive the benefits of standards-based SSO and API security integration without the cost and complexity of deploying a complete identity management (IdM) system. The PingFederate server integrates and coexists with existing home-grown and commercial IdM systems and applications, using these key features available separately from Ping Identity.

- **Integration Kits** – These tailored kits simplify integration with existing applications while minimizing impacts on existing infrastructure.
- **Token Translators** – These specialized plug-ins connect the STS with Web Service Providers and Clients to enable access to identity-enabled Web Services, which may require a range of different token types.

- **SaaS Connectors** – These plug-ins provide quick-connection templates and automated user provisioning and deprovisioning for selected SaaS providers, including Salesforce and Google Apps.
- **Cloud Identity Connectors** – These plug-ins allow cloud identity providers such as Facebook, Yahoo!, Google and Salesforce to authenticate and connect users to SSO-enabled applications.

Integration kits

PingFederate provides a suite of integration kits to complete the first- and last-mile integration with your existing IdM systems and web applications. PingFederate integration kits are available for download from the Ping Identity [Downloads](#) website, take only minutes to install, and are configured from within the PingFederate administrative console.

Integration kits enable rapid session integration with both existing authentication services and target applications. In addition, PingFederate includes a Software Development Kit for creating custom integrations.

For more information, see [SSO integration kits and adapters](#) on page 67 in the “Key Concepts” chapter of the PingFederate *Administrator's manual*.

Token translators

Ping Identity offers special token processors (for an IdP) and token generators (for an SP) to enable the WS-Trust STS to validate and issue a variety of token types. These plug-ins, which supplement built-in SAML token processing and generation, are designed to handle local identity tokens required in a variety of security contexts.

For more information, see [Token processors and generators](#) on page 59 in the “Key Concepts” chapter of the PingFederate *Administrator's manual*.

SaaS connectors

SaaS connectors offer a streamlined approach for browser-based SSO to selected SaaS providers—including automatic user provisioning and deprovisioning (see [Outbound provisioning for IdPs](#) on page 81 in the “Key Concepts” chapter of the PingFederate *Administrator's manual*). The Connector packages (available separately) include quick-connection templates, which automatically configure endpoints and other connection information for each provider.

Cloud identity connectors

Ping Identity offers social identity integration with social networking sites. The OpenID cloud-identity connector leverages OpenID 2.0 social networking providers (including Google and Yahoo!) for registration and access to cloud-based applications. Connectors for Twitter, LinkedIn, Windows Live, and Facebook leverage user logins for registration and access to cloud-based applications.

About PingOne

[PingOne](#) is Ping Identity's multi-tenant, identity-as-a-service (IDaaS) solution. PingOne® enables browser-based SSO and user provisioning for Identity Providers, and provides application providers with a rapid-deployment SSO capability. PingOne can be used together with PingFederate® to provide a powerful solution combining the benefits of an on-premise deployment with the flexibility of a cloud solution.

For more information on PingOne, please visit pingone.com.

Installation

PingFederate® is packaged as a stand-alone server based on J2EE application server technology. A new installation involves the following tasks:

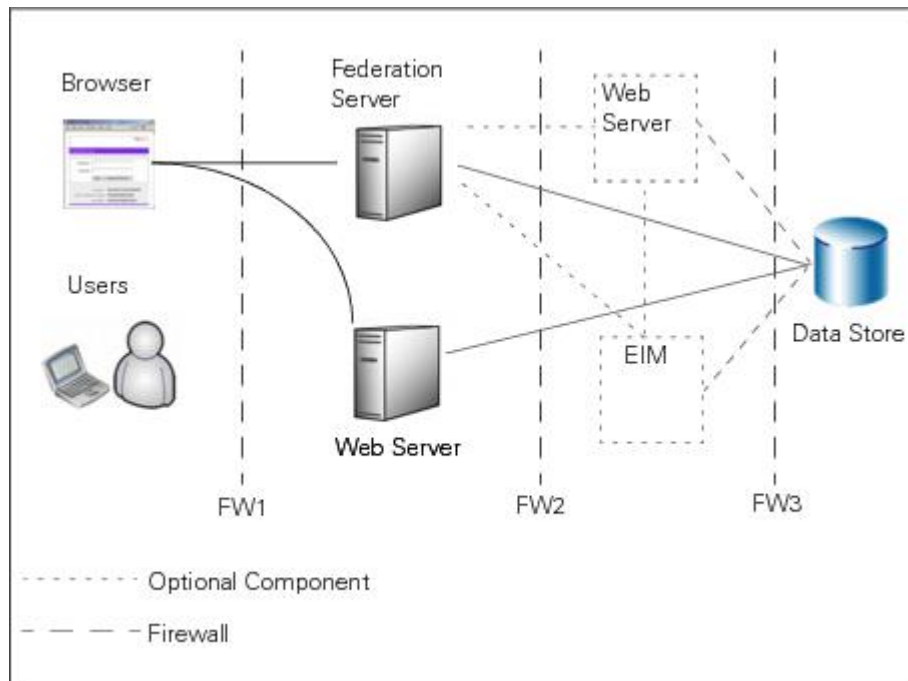
- [Deployment options](#) on page 13
- [System requirements](#) on page 14
- [Port requirements](#) on page 17
- [Install Java](#) on page 19

- [Installation options](#) on page 19
- [Initial setup wizard](#) on page 22

Deployment options


There are many options for deploying PingFederate® in your network environment, depending on your needs and infrastructure capabilities.

For example, you can choose a stand-alone or proxy configuration. The following diagram illustrates PingFederate installed in a demilitarized zone (DMZ):



In this configuration, the users access PingFederate via a web application server, an enterprise identity management (EIM) system, or both. PingFederate may, in turn, retrieve information from a data store 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 the proxy is accessed by users and web browsers. The proxy, in turn, communicates with PingFederate to request SSO.




System requirements

PingFederate® is certified as compatible for deployment and configuration with the minimum system specifications defined below.

Software requirements

Ping Identity has qualified the following configurations and certified that they are compatible with the product. Variations of these platforms (for example, differences in operating system version or service pack) are supported up until the point at which an issue is suspected as being caused by the platform or other required software.

Operating systems

 **Note:** PingFederate has been tested with default configurations of operating-system components. If your organization has customized implementations or has installed third-party plug-ins, deployment of the PingFederate server may be affected.

- Canonical Ubuntu 16.04 LTS
- Microsoft Windows Server 2012 Standard
- Microsoft Windows Server 2012 R2 Datacenter
- Microsoft Windows Server 2016
- Oracle Enterprise Linux 6.8 (Red Hat compatible kernel)
- Oracle Enterprise Linux 7.3 (Red Hat compatible kernel)
- Oracle Solaris 10
- Red Hat Enterprise Linux ES 6.8

- Red Hat Enterprise Linux ES 7.3
- SUSE Linux Enterprise 11 SP4
- SUSE Linux Enterprise 12 SP2

Docker support

- Docker version: 1.12.1
- Host operating system: Canonical Ubuntu 16.04 LTS
- Kernel: 4.4.0-38-generic

Virtualization

Although Ping Identity does not qualify or recommend any specific virtual-machine (VM) or container products other than those listed above, PingFederate has been shown to 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

Oracle Java SE Runtime Environment (Server JRE) 8 update 131

Browsers

Runtime server

- Chrome
- Edge
- Firefox
- Internet Explorer 11
- Safari
- Android 7 (Chrome)
- iOS 10 (Safari)

Administrative console and API

- Chrome
- Firefox
- Internet Explorer 11

TLS protocols

Runtime server

- TLS 1.2, 1.1, and 1.0

Administrative console and API

- TLS 1.2 and 1.1

Data store integration



Note: The following lists include data stores tested against specific PingFederate use cases. However, comparable data stores for each feature may potentially include any LDAP v3-compliant directory service or JDBC-compatible database.

For user-attribute lookup

- PingDirectory™ 6.0
- Microsoft Active Directory (2008 R2 and 2012)

- Oracle Directory Server Enterprise Edition 11g
- Microsoft SQL Server (2014 and 2016)
- Oracle Database (11g R2 and 12c)
- Oracle MySQL 5.7
- PostgreSQL 9.6.1

For SaaS or SCIM outbound provisioning

- Provisioning channel data source
 - PingDirectory 6.0
 - Microsoft Active Directory (2008 R2 and 2012)
 - Oracle Directory Server Enterprise Edition 11g
- Provisioning internal data store
 - Microsoft SQL Server (2014 and 2016)
 - Oracle Database (11g R2 and 12c)
 - Oracle MySQL 5.7
 - PostgreSQL 9.6.1

For SCIM inbound provisioning

- Microsoft Active Directory (2008 R2 and 2012)

For just-in-time (JIT) inbound provisioning (external target database)

- Microsoft SQL Server (2014 and 2016)

For account linking

- PingDirectory 6.0
- Microsoft Active Directory (2008 R2 and 2012)
- Oracle Directory Server Enterprise Edition 11g
- Microsoft SQL Server (2014 and 2016)
- Oracle Database (11g R2 and 12c)
- Oracle MySQL 5.7
- PostgreSQL 9.6.1

For OAuth client configuration

- Microsoft SQL Server (2014 and 2016)
- Oracle Database (11g R2 and 12c)
- Oracle MySQL 5.7
- PostgreSQL 9.6.1

For OAuth persistent grants

- PingDirectory 6.0
- Microsoft SQL Server (2014 and 2016)
- Oracle Database (11g R2 and 12c)
- Oracle MySQL 5.7
- PostgreSQL 9.6.1
- Microsoft Active Directory 2012
- Oracle Directory Server Enterprise Edition 11g
- Custom implementation through the PingFederate SDK



Note: PingFederate has been tested with vendor-specific JDBC 4.1 drivers. To obtain your database driver JAR file, contact your database vendor.

Hardware security module (optional)

- SafeNet Luna SA version 6.2.1
Client Driver Version: 5.3
- Thales nShield Connect version 2.51.10
Client Driver Version: 11.70

Hardware requirements



Note: Although it is possible to run PingFederate on less powerful hardware, the following guidelines accommodate disk space for default logging and auditing profiles and CPU resources for a moderate level of concurrent request processing.

Minimum hardware requirements

- Intel Pentium 4, 1.8 GHz processor
- 1 GB of RAM
- 1 GB of available hard drive space

Minimum hardware recommendations

- Multi-core Intel Xeon processor or higher
4 CPU/Cores recommended
- Multi-core SPARC processor (Solaris)
4 CPU/Cores recommended
- 4 GB of RAM
1.5 GB available to PingFederate
- 1 GB of available hard drive space

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.



Note: *Direction* refers to the direction of the initial requests relative to PingFederate. *Inbound* refers to requests received by PingFederate from external components. *Outbound* refers to requests sent by PingFederate to external components.


Service (Type of Traffic)	Protocol, Direction, Transport, Default Port	Source	Destination	Description
Administrative console	HTTPS, inbound, TCP, 9999	Administrator browser, administrative API REST calls, web service calls to the Connection Management Service	Administrative node ¹	Used for incoming requests to the administrative console. Configurable in the <code>run.properties</code> file.

¹ In a standalone environment, your PingFederate server is both the administrative node and the runtime engine node. In a clustered environment, you configure one of your Pingfederate as the sole administrative node in the cluster and the rest of the PingFederate servers as runtime engine nodes.

Service (Type of Traffic)	Protocol, Direction, Transport, Default Port	Source	Destination	Description
		Applicable to the console node in a clustered PingFederate environment		
Runtime engine	HTTPS, inbound, TCP, 9031	Client browser; mobile devices; web service calls to the SSO Directory Service, the OAuth Client Management Service, and the OAuth Access Grant Management Service; Session Revocation API REST calls	Runtime engine nodes	Used for incoming requests to the runtime engine. Configurable in the <code>run.properties</code> file. ²
		Applicable to all runtime engine nodes in a clustered PingFederate environment		
Cluster traffic (TCP)	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 <code>run.properties</code> file.
Cluster traffic (TCP)	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 <code>run.properties</code> file.
Cluster traffic (TCP, optional)	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 (UDP)	JGroups, inbound, UDP, 7601	PingFederate peer servers in a clustered	Administrative node and runtime engine nodes	Used for communications between engine nodes in a cluster when the transport mode for cluster


² The `pf.secondary.https.port`, if activated in the `run.properties` file, needs to be open as well.

Service (Type of Traffic)	Protocol, Direction, Transport, Default Port	Source	Destination	Description
		PingFederate environment		traffic is set to UDP. By default, the transport mode is TCP. Configurable in the <code>run.properties</code> file.
PingOne® integration (optional)	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, monitoring of PingFederate from PingOne, authenticating end users against the PingOne directory.
Active Directory domains/ Kerberos realms (optional)	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.


 **Note:** Depending on the integration kits deployed and the connecting third-party systems (such as email server or SMS service provider), you may need to open additional ports.

Install Java

You must install the Oracle Java SE Runtime Environment (Server JRE) before running PingFederate®, see [System requirements](#) for more information.

 **Tip:** Due to import control restrictions, the standard Server JRE distribution supports strong but not unlimited encryption. Stronger encryption is optional in several PingFederate and plug-in configurations. To use the strongest encryption, when permissible, after installing the Server JRE, download and install the appropriate version of “Java Cryptography Extension (JCE) Unlimited Strength Jurisdiction Policy Files” from the [Oracle Downloads](#) website (www.oracle.com/technetwork/java/javase/downloads/index.html).

1. Download and install the Server JRE from the [Oracle Downloads](#).
2. Set the `JAVA_HOME` environment variable to the Server JRE installation directory path and add its `bin` directory to the `PATH` environment variable.


 **Note:** If you intend to run PingFederate as a service, you must set the `JAVA_HOME` variable and modify `PATH` variable at the system level; otherwise, you have the options to set the variables at either the system or user level.

Installation options

You install PingFederate® by extracting the distribution ZIP file or running a platform-specific installer (for Microsoft Windows Server or Red Hat Enterprise Linux).

Install PingFederate on Windows

1. Request a license key via the Ping Identity [licensing](#) website.
2. Ensure you are logged on to your system with appropriate privileges to install and run an application.
3. Verify that the Server JRE is installed and the required environment variables are set correctly (see [Install Java](#) on page 19).

 **Note:** If you intend to use the PingFederate® installer for Windows or run PingFederate as a service, you must set the `JAVA_HOME` variable and modify `PATH` variable at the system level; otherwise, you have the options to set the variables at either the system or user level.

4. Install PingFederate via the platform-specific installer or the distribution ZIP file.

Installation medium	Steps
----------------------------	--------------

Platform-specific installer

Download and run the PingFederate installer for Windows.

PingFederate is configured to run as a service and started automatically at the end of the installation process.



Note: The PingFederate installer for Windows is designed to install 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. Note that you must manually configure various port settings in the `<pf_install>/pingfederate/bin/run.properties` file (for each instance) to avoid any port conflicts.

Distribution ZIP file

Download and extract the distribution ZIP file into an installation directory.



Note: 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, in certain server-clustering scenarios), either install each instance in a different location or rename the parent directory to install a parallel file structure in the same location.

The installation directory is referred as `<pf_install>`, where the `pingfederate` directory can be found; for example, `<pf_install>/pingfederate`.

5. If you have installed PingFederate by 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...
```



Tip: To configure PingFederate to run as a service, follow the steps in [Install the PingFederate service on Windows manually](#) on page 25.



Note: 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 51.

Install PingFederate on Red Hat Enterprise Linux

1. Request a license key via the Ping Identity [licensing](#) website.
2. Ensure you are logged on to your system with appropriate privileges to install and run an application.



Note: You must install and run PingFederate® under a local user account.

3. Verify that the Server JRE is installed and the required environment variables are set correctly (see [Install Java](#) on page 19).
4. Install PingFederate via the platform-specific installer or the distribution ZIP file.

Installation media	Steps
---------------------------	--------------


Platform-specific installer

Download and run the PingFederate install script.

PingFederate is configured to run as a service and started automatically at the end of the installation process.

Distribution ZIP file

Download and extract the distribution ZIP file into an installation directory.

 **Note:** 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, in certain server-clustering scenarios), either install each instance in a different location or rename the parent directory to install a parallel file structure in the same location.


The installation directory is referred as `<pf_install>`, where the `pingfederate` directory can be found; for example, `<pf_install>/pingfederate`.


5. If you have installed PingFederate by the distribution ZIP file, start PingFederate manually by running the following script:

```
<pf_install>/pingfederate/bin/run.sh
```

Wait for the script to finish—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 [Install the PingFederate service on Linux manually](#) on page 25.

 **Note:** 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 51.


Install PingFederate on UNIX/Linux

Refer to [System requirements](#) on page 14 for a list of qualified UNIX and Linux operating systems.

1. Request a license key via the Ping Identity [licensing](#) website.
2. Ensure you are logged on to your system with appropriate privileges to install and run an application.

 **Note:** You must install and run PingFederate® under a local user account.

3. Verify that the Server JRE is installed and the required environment variables are set correctly (see [Install Java](#) on page 19).
4. Download and extract the distribution ZIP file into an installation directory.

 **Note:** 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, in certain server-clustering scenarios), either install each instance in a different location or rename the parent directory to install a parallel file structure in the same location.


The installation directory is referred as `<pf_install>`, where the `pingfederate` directory can be found; for example, `<pf_install>/pingfederate`.


5. Start PingFederate manually by running the following script:

```
<pf_install>/pingfederate/bin/run.sh
```

Wait for the script to finish—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 [Install the PingFederate service on Linux manually](#) on page 25.

 **Note:** 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 51.

Open the PingFederate administrative console

The administrator's user interface for PingFederate®, the *administrative console*, is built around a system of wizard-like control screens, in which you configure various settings and components to support your federation use cases.

1. Start a web browser; for example, Chrome, Firefox, or Internet Explorer 11.

2. Browse to the following URL:

```
https://<DNS_NAME>:9999/pingfederate/app
```

where <DNS_NAME> is the fully qualified name of the machine running the PingFederate server.

Initial setup wizard

The first time you run the PingFederate® administrative console, the initial setup wizard guides you through the process of configuring your identity federation settings and optionally connecting PingFederate to PingOne® to deploy a powerful on-premise and cloud-based hybrid solution. The tasks include:

- [Connect PingFederate to PingOne](#) on page 22 (optional)
- [Review or import your license](#) on page 22
- [Enter the basic information](#) on page 23
- [Select your federation roles](#) on page 23
- [Configure identity provider settings](#) on page 23 (optional)
- [Create an administrator account](#) on page 25
- [Review the initial configuration](#) on page 25

Connect PingFederate to PingOne

PingOne is Ping Identity's multi-tenant, identity-as-a-service (IDaaS) solution. PingOne® enables browser-based SSO and user provisioning for Identity Providers, and provides application providers with a rapid-deployment SSO capability. PingOne can be used together with PingFederate® to provide a powerful solution combining the benefits of an on-premise deployment with the flexibility of a cloud solution.

When you select to connect to PingOne, a *managed* PingOne SP connection is created for you. In this managed SP connection, PingFederate automatically uploads any metadata changes to PingOne. For example, if you update your **Base URL** under Server Settings or extend the attribute contract to PingOne, those changes will be propagated to PingOne without additional effort. In addition, PingFederate will automatically download new signing certificates from PingOne and update the connection accordingly.

1. Click **Sign on to PingOne to get your activation key**.
2. Sign on using your PingOne admin portal credentials.
3. Copy the **Activation Key** value.
4. Click **Finished**.
5. Close the browser tab and go back to the PingFederate administrative console.
6. In the **PingOne Account** screen, paste your activation key.
7. Click **Next**.

If you prefer to setup PingFederate without PingOne for now, click **Next** to continue. When you are ready to connect PingFederate to PingOne, go to the **Server Configuration** menu and click **Connect to PingOne**.

Review or import your license

- If you have selected to connect PingFederate® to PingOne®, the initial setup wizard automatically downloads a 30-day trial license after validating your activation key.

If you wish to import your license file, you may do so in the **License** screen.

Alternatively, you can import your license file in the **Server Configuration > License Management** screen *before* the trial license expires.

- If you have opted to setup PingFederate without PingOne, import your license in the **License** screen.




Note: If you do not have a PingFederate license yet, request a license key via the Ping Identity [licensing](#) website or contact sales@pingidentity.com.

Click **Next** to continue.


Enter the basic information

In the **Basic Information** screen, enter your federation information.

1. Verify your **Base URL**. Update as needed.

 **Tip:** 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® server.

2. Specify your **Entity ID**.


 **Note:** If you have selected to connect PingFederate to PingOne®, the **Entity ID** field is pre-populated for you based on your PingOne setup.

3. Click **Next**.

Select your federation roles

In the **Enable Roles** screen, select the roles of your PingFederate® server.

1. Select at least one role for your PingFederate server.

 **Note:** If you have selected to connect PingFederate to PingOne®, the **Identity Provider** role is selected for your convenience. You may select additional roles as needed.


2. Click **Next**.

Configure identity provider settings

The **Identity Provider Configuration** screen appears when the **Identity Provider** role is selected. Use this screen to:

- Connect your Active Directory as an LDAP data store
- Create adapters and authentication selectors to authenticate end users via the Kerberos protocol or a login form based on end-user browsers and your network topology
- Enable SSO from PingOne® to the PingFederate® administrative console or create a local administrative account to access the console

To continue with the initial setup wizard, click **Begin** or **Connect to Active Directory**.


 **Note:** The **Identity Provider Configuration** screen is also the second step in the *Connect to PingOne* wizard from the **Server Configuration** menu; this is the use case where you decided not to connect PingFederate to PingOne in the past but would like to do so now.

Connect your Active Directory

As an identity provider, you often need to supply additional information about your users, such as their first and last names and email addresses, to single sign-on to the SPs including PingOne®. If Active Directory (AD) is your user repository, use the **Connection** screen to establish a secure connection to your AD LDAP server.

Based on the information provided, the initial setup wizard also creates an LDAP Username Password Credential Validator instance and an HTML Form Adapter instance automatically for you, such that your users can authenticate through a login form using their AD credentials.

1. Enter the hostname and the access credentials.
2. Specify a **Search Base**. This is the starting point in your AD where PingFederate® looks for users and groups.
3. Modify the pre-populated **Search Filter** value as needed.
4. Click **Next**.

 **Note:** PingFederate tries to establish a secure connection to your AD via LDAPS.

If your AD LDAP server does not support LDAPS, the **Unsecure Connection** screen appears. If you want to continue without a secure connection, click **Next**.

If the subject of the certificate presented by your AD LDAP server does not match the **Hostname** value, the **Unsecure Connection** screen appears. Click **Previous** to update the **Hostname** field. If you want to continue without a secure connection, click **Next**.

If the certificate presented by your AD LDAP server is not trusted by PingFederate, the **Certificate Error** screen appears. Import the certificate used by your AD LDAP server to establish a secure connection or select the **I want to complete an unsecure connection** check box to continue without a secure connection, and then click **Next**.

Configure Kerberos authentication

PingFederate® is also capable of authenticating users using Active Directory credential tokens (specifically Kerberos service tickets), providing Windows users a seamless single sign-on experience.



Note: If you decide not to enable Kerberos authentication, users will authenticate through the HTML Form Adapter that was automatically created in the previous screen where you connected PingFederate to your Active Directory.

1. Select the **Configure Kerberos Authentication** check box.
2. Enter the realm name, the Kerberos service account and its password.



Important: If you have not created or configured a service account for Kerberos authentication, see [Configure the Active Directory environment](#) on page 204 for additional steps. You must have Domain Administrator permissions to make the required changes.

3. Optional: Enter one or more **KDC Hostnames**. If unspecified, PingFederate uses a DNS query to find a list of KDCs.
4. Optional: Click **Test** to verify the connectivity to your KDCs from the administrative console.

When a connection to any of the KDCs is successful, the message `Test was successful` appears. Otherwise, the test returns error messages near the top of the screen.

Note that the test stops at the first successful result, so all KDCs are not necessarily verified. Also, connectivity may be subsequently affected in different deployment scenarios, including for engine server nodes running in a clustered environment.

5. Enter one or more **Internal IP Ranges** in CIDR notation to indicate the boundaries of your network.

End users outside of your network will authenticate through the HTML Form Adapter created in the previous screen.



Note: End users using mobile clients, such as iPhone and Android mobile phones, will always authenticate through the HTML Form Adapter that was automatically created in the previous screen where you connected PingFederate to your Active Directory.

6. Click **Next**.



Important: You also need to configure the end-user browsers for seamless Kerberos authentication. For more information, see [Configure end-user browsers](#) on page 429.

Enable provisioning to PingOne

If you have selected to connect PingFederate® to PingOne®, the **Provisioning** screen appears and the **Configure Provisioning** check box is selected for your convenience.

This capability gives you the flexibility to provision users from PingOne to SaaS vendors when adding cloud applications later in the PingOne admin portal (see [Add an Application from the Application Catalog](#) in the PingOne *Employee SSO Administration Guide*).

Users and group provisioning also allows you to configure user access to cloud applications in the PingOne admin portal based on groups and membership information without waiting for end users to sign on (see [Manage Users by Group](#) in the PingOne *Employee SSO Administration Guide*).

1. Specify the **Group DN** where PingFederate should look for member users (under the **Search Base** previously defined in the Connection screen) to provision to PingOne.



Note: Groups under the **Search Base** are also provisioned to PingOne automatically.

2. Optional: Select the **Nested** check box to if you want PingFederate to provision users through nested group membership.

3. Click Next.

Review your identity provider configuration

- In the **Summary** screen, review your configuration.



Note: If you need to make any changes, click the heading over the information you want to edit. When you are finished, click **Next** multiple times until you are back to this **Summary** screen.

Click **Done** to continue.

Create an administrator account

Use the **Administrator Account** screen to create an administrative login.



Note: If you have selected to connect PingFederate® to PingOne®, this screen does not appear because the option to SSO to the PingFederate administrative console from PingOne has been enabled for you. As needed, you can manage the SSO option in the **Server Configuration > Server Settings > PingOne Settings** screen after the initial setup.

Click **Next** to continue.

Review the initial configuration

In the **Confirmation** screen, review your configuration.



Note: If you wish to edit any settings, select the respective screen. If you wish to start over, restart your browser.

In the **Complete** screen, click **Next** to apply the configuration, and then click **Done** to exit to the Main menu.

Install PingFederate as a service

You can set up PingFederate® to run in the background as a service on either Windows or Linux.



Tip: If you install PingFederate using one of the platform-specific installers, PingFederate has already been configured to run as a service and started automatically at the end of the installation process.



Important: When you start the server manually, you must also run the startup script under the same user account that the service uses.

Install the PingFederate service on Windows manually

If you have installed PingFederate® using the installer for Windows, PingFederate has already been configured to run as a service and started automatically at the end of the installation process. That is, you do not need to install the PingFederate service manually.

1. Install PingFederate.



Note: Ensure JAVA_HOME and PATH are set as system variables (see [Install Java](#) on page 19).

2. Ensure you are logged on with full Administrator privileges.

3. Start PowerShell or Command Prompt as an Administrator.

4. In PowerShell or Command Prompt, run the <pf_install>\pingfederate\sbin\win-x86-64\install-service.bat file to install the service.

5. Open the **Control Panel > Administrative Tools > Services** management console.


6. Right-click on the **PingFederate** service and select **Start**.

Similar to other services, the **PingFederate** service is installed and configured to start automatically on reboot.

Install the PingFederate service on Linux manually

If you have installed PingFederate® using the install script for Red Hat Enterprise Linux, PingFederate has already been configured to run as a service and started automatically at the end of the installation process. That is, you do not need to install the PingFederate service manually.

To run PingFederate as a service on Linux, you must place a script in the system initialization directory.

 **Note:** This topic uses `/etc/rc.d/init.d/` as the system initialization directory. Modify accordingly if the system initialization directory resides elsewhere on your Linux server.

1. Install PingFederate.
2. Log on as root.
3. Create a new user account for the service.


For this procedure, the variable `<pf_user>` is used to refer to this account.

 **Note:** Ensure the environment variable `JAVA_HOME` is set and the `PATH` variable updated for `<pf_user>` (see [Install Java](#) on page 19).

4. Change the PingFederate installation directory (`<pf_install>`) ownership and modify the read and write permissions:

```
chown -R <pf_user> <pf_install>
chmod -R 775 <pf_install>
```

5. Place the code below into a file called `<pf_user>` in the `/etc/rc.d/init.d/` directory.

 **Note:** Replace instances of `<pf_user>` and `<pf_install>` in the script below, and in the commands that follow, with their respective values.

Line breaks (as indicated by the backslash character, `\`) are inserted to the `su` command for readability. Remove them from the `su` command when creating the shell script.

```
#!/bin/sh

start(){
    echo "starting PingFederate.."
    su - <pf_user> \
    -c '<pf_install>/pingfederate/sbin/pingfederate-run.sh \
    > /dev/null 2> /dev/null'
}

stop(){
    echo "stopping PingFederate.."
    su - <pf_user> \
    -c '<pf_install>/pingfederate/sbin/\
    pingfederate-shutdown.sh'
}

restart(){
    stop
    # padding time to stop before restart
    sleep 60
    # To protect against any services that are not stopped,
    # uncomment the following command.
    # (Warning: this kills all Java instances running as
    # <pf_user>.)
    # su - <pf_user> -c 'killall java'
    start
}

case "$1" in
    start)
        start
        ;;
    stop)
        stop
        ;;
    restart)
        restart
        ;;
esac
```

```

*)
echo "Usage: <pf_user> {start|stop|restart}"
exit 1
esac
exit 0

```

6. Create symbolic links using commands listed below.

The links specify the order in which PingFederate starts and stops.

```

ln -s /etc/rc.d/init.d/<pf_user> /etc/rc3.d/S84<pf_user>
ln -s /etc/rc.d/init.d/<pf_user> /etc/rc5.d/S84<pf_user>
ln -s /etc/rc.d/init.d/<pf_user> /etc/rc4.d/S84<pf_user>
ln -s /etc/rc.d/init.d/<pf_user> /etc/rc6.d/K15<pf_user>
ln -s /etc/rc.d/init.d/<pf_user> /etc/rc0.d/K15<pf_user>
ln -s /etc/rc.d/init.d/<pf_user> /etc/rc1.d/K15<pf_user>
ln -s /etc/rc.d/init.d/<pf_user> /etc/rc2.d/K15<pf_user>

```

7. Make the script executable (as root):

```
chmod 755 /etc/rc.d/init.d/<pf_user>
```

8. Test the script using the `service` command.

```

service <pf_user> start
service <pf_user> stop
service <pf_user> restart

```



Note: Some Linux systems may require a reboot before the script can be tested.

Uninstall PingFederate

Uninstalling PingFederate® involves removing the PingFederate service (if it was previously installed) and the installation directory (`pf_install`).

Uninstall PingFederate from a Windows server

1. Ensure you are logged on to your system with appropriate privileges to uninstall an application.
2. Verify the installation medium in **Control Panel > Uninstall a Program**.

A **PingFederate** entry indicates that PingFederate® was previously installed using the PingFederate installer for Windows; otherwise, it was installed using a distribution ZIP file.

3. Uninstall PingFederate.

Installation medium	Steps
PingFederate installer for Windows	<ol style="list-style-type: none"> 1. (Optional) Make a backup copy of the PingFederate installation directory (<code>pf_install</code>). 2. Use Control Panel > Uninstall a Program to uninstall PingFederate, which removes the PingFederate service and the installation directory.
Distribution ZIP file	<ol style="list-style-type: none"> 1. Open the Control Panel > Administrative Tools > Services management console. 2. Right-click on the PingFederate service (if found) and select Stop, and then run <code>uninstall-service.bat</code> from the <code>pf_install\pingfederate\sbin</code> subdirectory that corresponds to your platform processor. 3. (Optional) Remove the PingFederate installation directory (<code>pf_install</code>).


Uninstall PingFederate from a Linux server

1. Log on as the root user.
2. Stop the service with the command:


```
service <pf_user> stop
```

 where *<pf_user>* is the PingFederate® service user account.
3. Remove symbolic links:

```
rm /etc/rc3.d/S84<pf_user>
rm /etc/rc4.d/S84<pf_user>
rm /etc/rc5.d/S84<pf_user>
rm /etc/rc0.d/K15<pf_user>
rm /etc/rc1.d/K15<pf_user>
rm /etc/rc2.d/K15<pf_user>
rm /etc/rc6.d/K15<pf_user>
```

 **Note:** This topic uses `/etc/rc.d/init.d/` as the system initialization directory. Modify accordingly if the system initialization directory resides elsewhere on your Linux server.

4. Optional: Delete the script used to start and stop the service.
5. Optional: Remove the PingFederate installation directory (`<pf_install>`).

The administrative console

The administrator's user interface for PingFederate®, the *administrative console*, is built around a system of wizard-like control screens. To access the administrative console, launch a web browser and go to:

`https://<DNS_NAME>:9999/pingfederate/app`

where `<DNS_NAME>` is the fully qualified name of the machine running the PingFederate server.

Once logged on, the Main menu offers up to four menu choices:

- IdP Configuration
- SP Configuration
- OAuth Settings
- Server Configuration

Note that the menu choices and menu items varies with the federation roles PingFederate plays; for example, IdP, SP, SCIM provisioner, or OAuth Authorization Server (see *Choose roles and protocols* on page 151 in the “System Settings” chapter of the PingFederate *Administrator's manual*). Menu items also depend on the permissions assigned to the logged-on administrator (see *Account management* on page 108 in the “System Administration” chapter of the PingFederate *Administrator's manual*).

Tasks and steps

Each broad configuration area is broken down into a series of tasks. Each task consists of a sequence of steps. The tasks and steps appear in the top portion of the screen.




Figure 1: Sample tasks and steps

In this example, the primary task is managing one or more IdP adapter instances (**Manage IdP Adapter Instances**). The administrator is working on creating an adapter instance (**Create Adapter Instance**). The current step is about selecting the type the IdP adapter (**Type**). The subsequent steps, which the administrator has not yet reached, are grayed out.

The administrator console displays a summary screen at the end of a task, offering the opportunity to review and make changes as needed.

Some steps provide buttons that branch to dependent tasks with multiple steps. When the dependent tasks are complete, the administrative console brings back the originating tasks for the administrators to continue with the rest of the configuration.

For instance, when creating a connection to a partner, the administrator might need to create a new digital signing certificate, which is a common task with its own set of steps. The connection wizard provides a button to open the task of creating a new signing certificate. When the administrator completes the task, the administrative console brings the administrator back to the connection wizard to finish off the configuration of it.

Note that clicking **Cancel** on any screen discards all new unsaved entries or changes for all steps shown for the current task and returns you to the screen from which you accessed the task.

Console buttons

The navigational and control buttons at the bottom of the administrative console screen change depending on where you are in the configuration process. The following table describes the behavior of these buttons:

Button	Description
Save	Stores information for all steps completed for the current task or any changes made for the current step; returns to the screen from the which the task or step was accessed. This button is available only when the Save operation is valid within the current context.
Done	Marks as complete all steps for a current task, but does not save the configuration (because further tasks or steps are necessary); to save entries or changes, click Save (or continue the configuration until you see a Save button). When creating a new connection, click Save Draft (see below).
Save Draft	Stores a new connection configuration for all steps completed up to the current screen in the configuration flow. To return to the draft, click Manage All under SP or IdP Connections and then select the draft from the connection list.
Cancel	Returns to the screen from which the current task was accessed; discards any information newly entered or modified for all steps in the task.
Previous	Returns to the previous step (when applicable).
Next	Moves display forward to the next step (when applicable), if all required information is complete in the current step.



Caution: Do not use the browser's Back, Refresh, or Forward buttons. Instead, always use the navigation buttons, **Previous**, **Next**, or **Done**.

Supported standards

PingFederate® provides flexible, integrated support for the Security Assertion Markup Language (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.

This chapter describes:

- [Federation roles](#) on page 30
- [Terminology](#) on page 30
- [SAML 1.x profiles](#) on page 32
- [SAML 2.0 profiles](#) on page 34
- [WS-Federation](#) on page 44

- [About account linking](#) on page 45
- [Web Services standards](#) on page 46
- [OAuth 2.0 and PingFederate AS](#) on page 48
- [System for Cross-domain Identity Management \(SCIM\)](#) on page 50
- [Transport and message security](#) on page 51

Federation roles

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, however, 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.

With OAuth 2.0 and OpenID Connect 1.0 support, PingFederate can be configured as an authorization server (AS), an OpenID Provider (OP), and a Relying Party (RP). (Note that OP and RP are the synonyms for IdP and SP, respectively.)

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 that authentication.



Note: The SAML subject may be a person, a web application, or a web server. Since the subject is often a person, the term *user* is generally employed throughout our documentation.

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 whether (or 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 it can be employed as a stand-alone server.

Authorization server

An OAuth AS issues access tokens and refresh tokens to OAuth clients after the resource owner has fulfilled the authentication requirement.

OpenID Provider

An 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

The SAML specifications provide a system of building blocks and support components for achieving secure data exchange in an identity federation. These include:

- **Assertions**
- **Bindings**

- **Profiles**
- **Metadata**
- **Authentication Context**

Assertions

Assertions are XML documents sent from an IdP to an SP. Each assertion contains identifying information about a user who has initiated an SSO request.

Bindings

A SAML binding describes the way messages are exchanged using transport protocols. PingFederate® supports the following bindings:

HTTP POST

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

HTTP Artifact

Describes how to use an artifact to represent a SAML message. The artifact can be transported via an HTML form control or a query string in the URL.

HTTP Redirect (SAML 2.0)

Describes how SAML messages are transported using HTTP 302 status-code response messages.

SOAP (SAML 2.0)

Describes how SAML messages are to be transferred across the back channel (Simple Object Access Protocol).

Profiles

Profiles describe processes and message flows combining assertions, request/response message specifications, and bindings to achieve a specific desired functionality or use case. Because profiles define the application of the specifications and therefore play a large part in PingFederate, most of the rest of this chapter is devoted to them, starting with [SAML 1.x profiles](#) on page 32.

Metadata

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

Whether you are publishing or consuming 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 may want information surrounding how the user was originally authenticated by the IdP, in addition to the assertion itself. The SP may use 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 can create authentication-context declarations. Partners may 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: [saml-authn-context-2.0-os.pdf](#)). However, it is up to partners to decide if additional authentication context is required and if these classes supply an adequate description. For SAML 1.x, the authentication context (called *AuthenticationMethod*), if used, must be specified 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.

Alternatively, several PingFederate integration kits provide methods that can be used by the developer to insert authentication context from external IdP applications into the assertion. Conversely, the SP developer can call

methods for extracting authentication context from an assertion. Ultimately, it is up to the SP developer and application to create access control or other processing based on the context.

Browser-based SSO

Browser-based SSO includes SAML 1.x, 2.0, WS-Federation, and OpenID Connect and provides standards-based SSO, Single Logout, Attribute Query and XASP, and the WS-Federation Passive Requestor Profile for SP-initiated SSO.

SAML 1.x profiles

SAML 1.0 and 1.1 profiles provide for browser-based SSO, initiated by an IdP, using either the POST or artifact bindings.

In addition, the specifications provide for a non-normative SP-initiated scenario (called “destination-first”), which allows web developers to create applications that enable a user to initiate SSO from the SP site.

SSO—Browser-POST

In this scenario, 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 via HTTP POST.

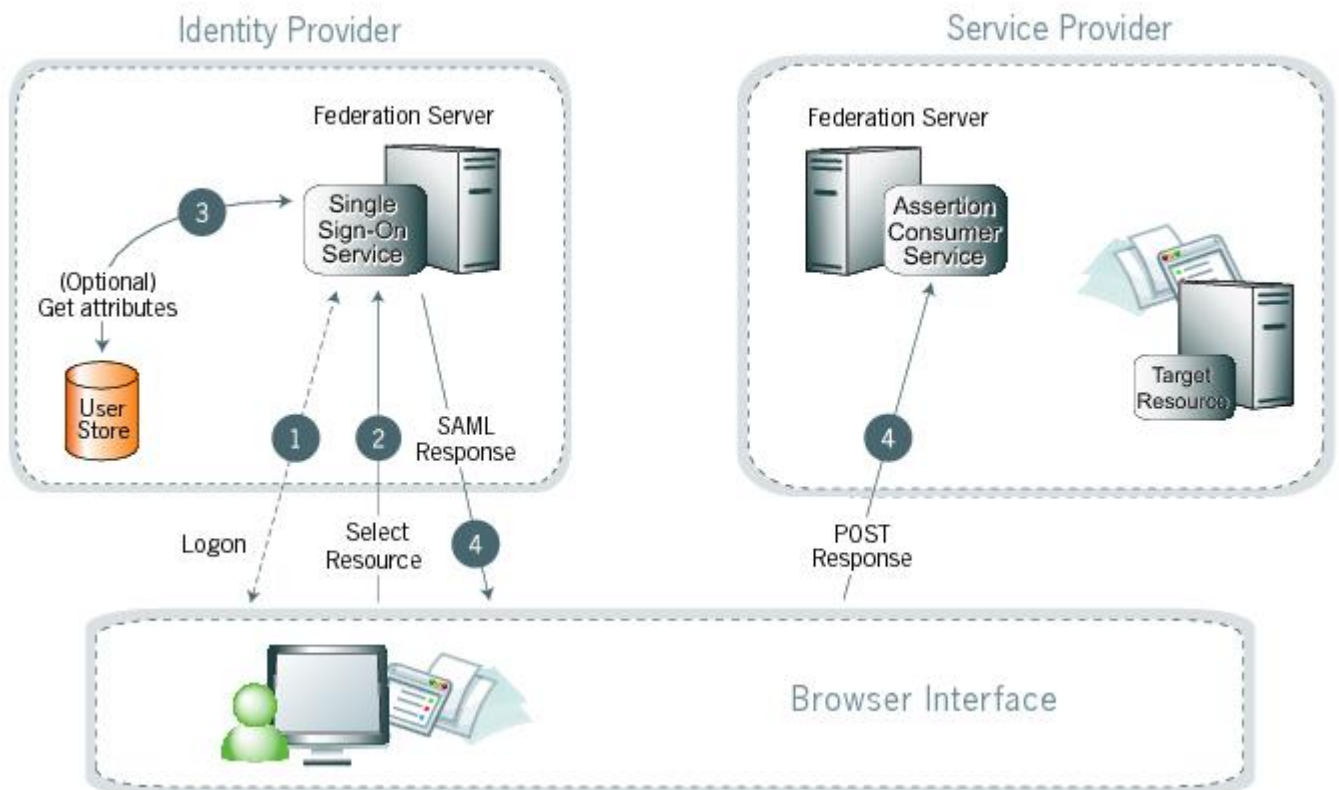



Figure 2: Browser/POST Profile

Processing steps

1. A user has logged 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. 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 that POST responses be digitally signed.

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

SSO—Browser-Artifact

In this scenario, the IdP sends a SAML artifact to the SP via either HTTP POST or a redirect (shown in diagram). The SP uses the artifact to obtain the associated SAML response from the IdP.

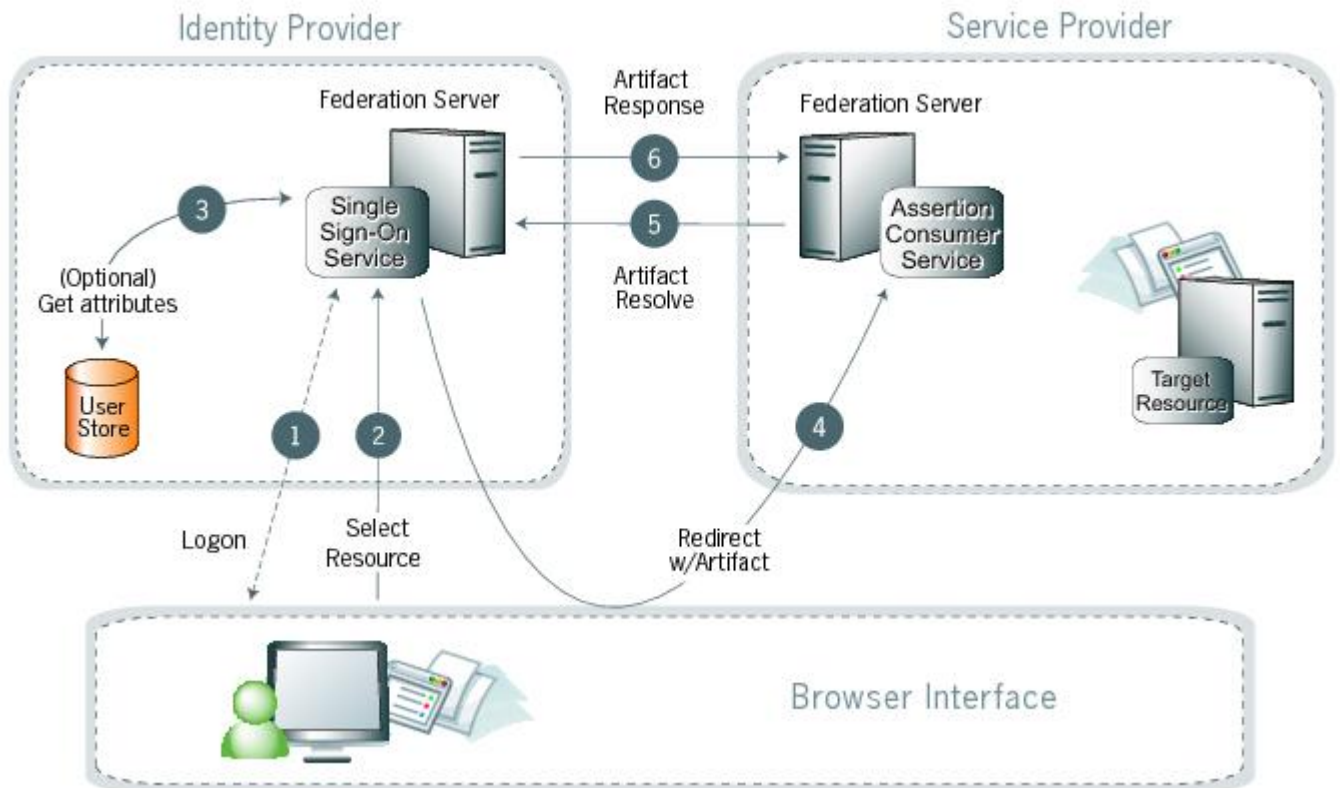


Figure 3: SSO:Browser/Artifact Profile

Processing steps

- A user has logged 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).
- The user clicks a link or otherwise requests access to a protected SP resource.
- Optionally, the IdP retrieves attributes from the user data store.
- 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).
- 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).
- The ARS sends a SAML artifact response message containing the previously generated assertion.
- (Not shown) If a valid assertion is received, the SP establishes a session and redirects the browser to the target resource.

SP-initiated ("destination-first") SSO

In an SP-initiated (a.k.a. "destination-first") transaction the user is connected 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 federation agreement, authentication is managed by the IdP. The SP sends an authentication request to the IdP.

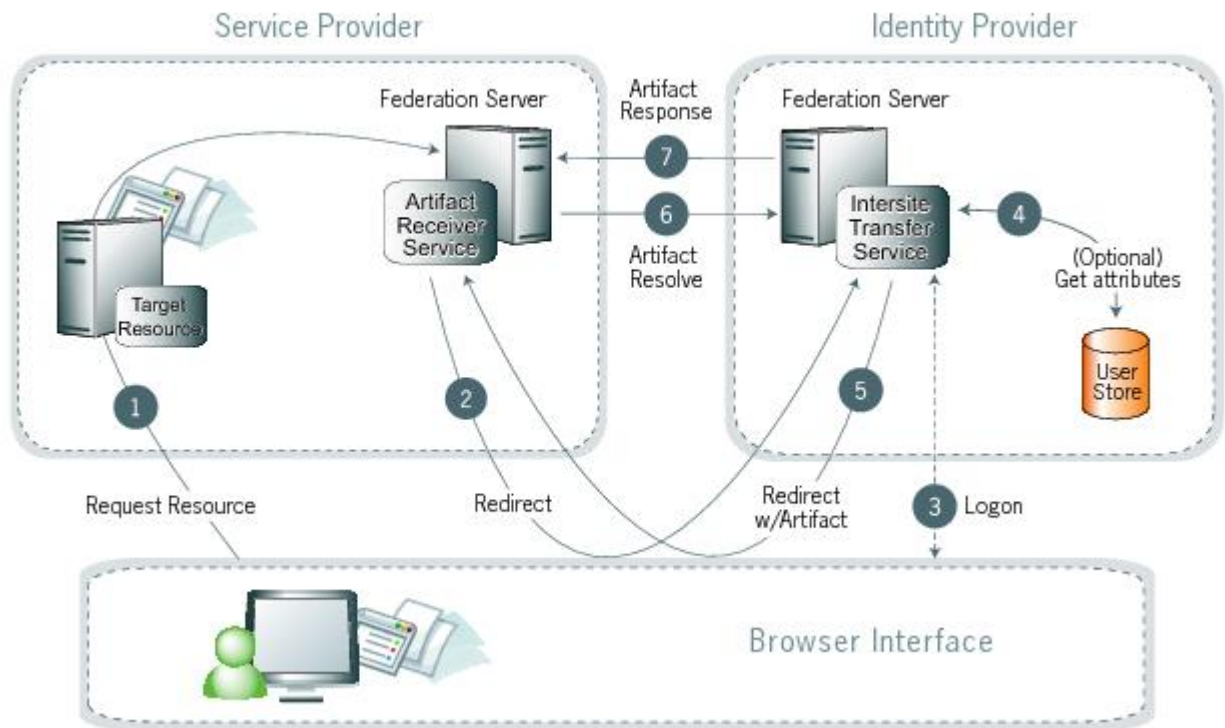


Figure 4: SP-Initiated SSO

Processing steps

1. The user requests access to a protected SP resource. The request is redirected to the federation server (for example, PingFederate®) to handle authentication.
2. The federation server sends a SAML request for authentication to the IdP's SSO service (also called the Intersite Transfer Service).
3. If the user is not already logged on to the IdP site or if re-authentication is required, the IdP asks for credentials (e.g., ID and password) and the user logs on.
4. Additional information about the user may be retrieved from the user data store for inclusion in the SAML response. (These attributes are predetermined as part of the federation agreement between the IdP and the SP—see [About attributes](#) on page 70 in the “Key Concepts” chapter of the PingFederate *Administrator's manual*.)
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 endpoint.
7. The IdP resolves the artifact and returns the corresponding SAML response with the SSO assertion.
8. (Not shown) If the assertion is valid, the SP establishes a session for the user and redirects the browser to the target resource.

SAML 2.0 profiles

PingFederate® supports these major profiles defined under the SAML 2.0 standard:

- [Single sign-on](#) on page 35
- [Single logout](#) on page 42
- [Attribute Query and XASP](#) on page 43
- [Standard IdP Discovery](#) on page 43

Single sign-on

SAML 2.0 substantially increases the number of possible SSO profile variations by fully enabling SP-initiated transactions. When SP- and IdP-initiated protocols are paired with transport binding specifications, the combinations result in eight practical SSO scenarios:

- [SP-initiated SSO—POST-POST](#) on page 35
- [SP-initiated SSO—Redirect-POST](#) on page 36
- [SP-initiated SSO—Artifact-POST](#) on page 37
- [SP-initiated SSO—POST-Artifact](#) on page 38
- [SP-initiated SSO—Artifact-Artifact](#) on page 40
- [SP-initiated SSO—Redirect-Artifact](#) on page 39
- [IdP-initiated SSO—POST](#) on page 41
- [IdP-initiated SSO—Artifact](#) on page 42

SP-initiated SSO—POST-POST

In this scenario a user attempts to access a protected resource directly on an SP website without being logged on. The user does not have an account on the SP site, but does have a federated account managed by a third-party IdP. The SP sends an authentication request to the IdP. Both the request and the returned SAML assertion are sent through the user's browser via HTTP POST.

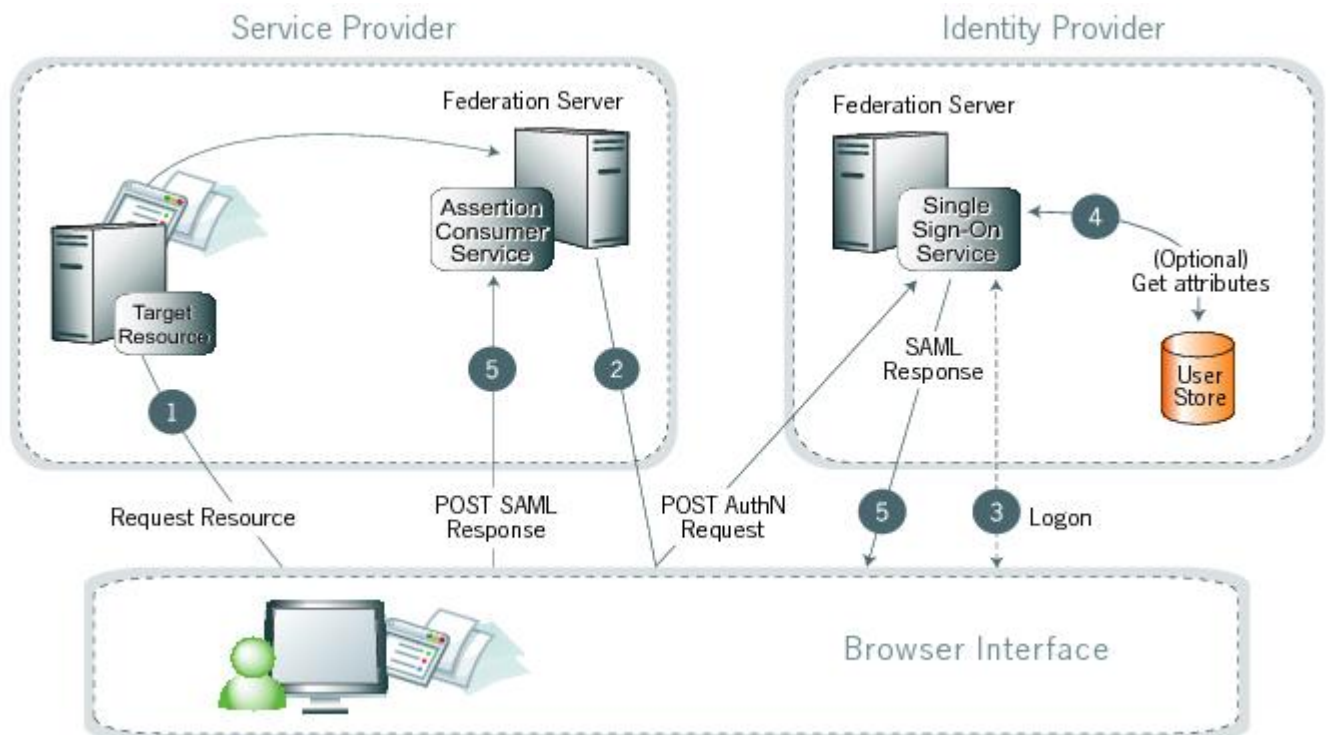


Figure 5: SP-Initiated SSO: POST/POST

Processing steps

1. The user requests access to a protected SP resource. The request is redirected 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 is automatically posted to the IdP's SSO service.
3. If the user is not already logged on to the IdP site or if re-authentication is required, the IdP asks for credentials (e.g., ID and password) and the user logs on.

4. Additional information about the user may be retrieved from the user data store for inclusion in the SAML response. (These attributes are predetermined as part of the federation agreement between the IdP and the SP—see [About attributes](#) on page 70 in the “Key Concepts” chapter of the PingFederate *Administrator's manual*.)
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 that POST responses be digitally signed.

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 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.

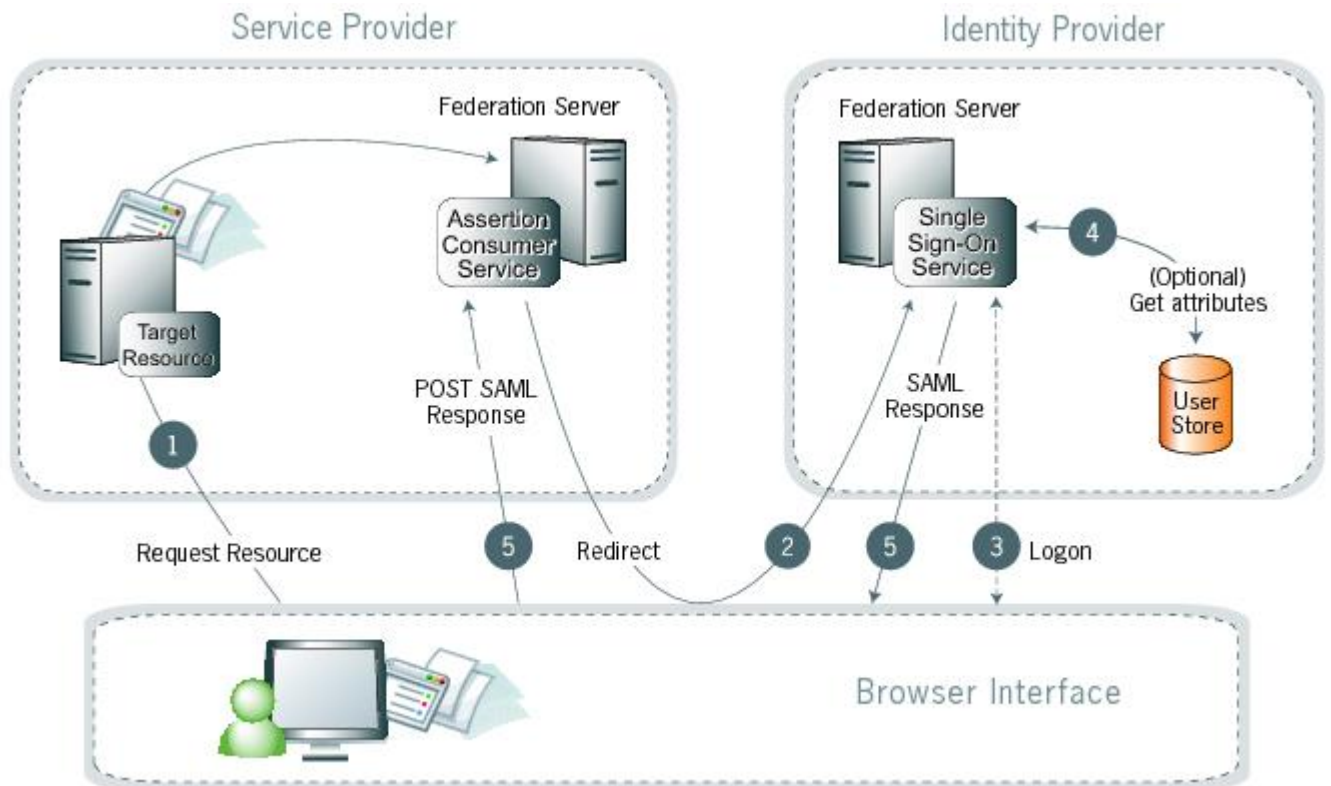



Figure 6: SP-Initiated SSO: Redirect/POST

Processing steps

1. A user requests access to a protected SP resource. The user is not logged on to the site. The request is redirected 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 SSO service.
3. If the user is not already logged on to the IdP site or if re-authentication is required, the IdP asks for credentials (e.g., ID and password) and the user logs on.
4. Additional information about the user may be retrieved from the user data store for inclusion in the SAML response. (These attributes are predetermined as part of the federation agreement between the IdP and the SP—see [About attributes](#) on page 70 in the “Key Concepts” chapter of the PingFederate *Administrator's manual*.)
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 that POST responses be digitally signed.

- (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—Artifact-POST

In this scenario, the SP sends a SAML artifact to the IdP via 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.

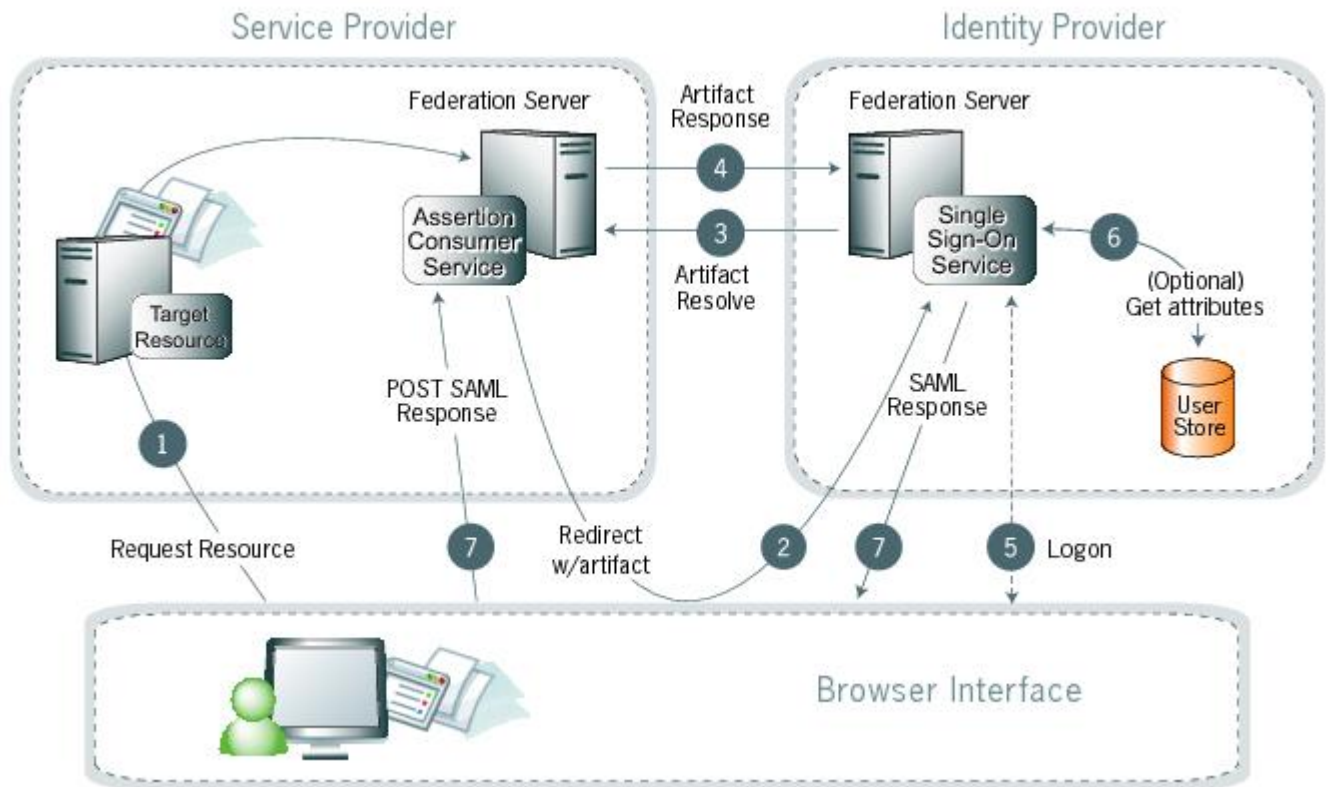





Figure 7: SP-Initiated SSO: Artifact/POST

Processing steps

- 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.
- 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 SSO service.
 -  **Note:** The artifact contains the source ID of the SP's artifact resolution service and a reference to the authentication.
- 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 Artifact Resolution Service (ARS).
 -  **Note:** The SP and IdP's source IDs and remote artifact resolution services are mapped according to the federation agreement made prior to this action.
- The SP's ARS returns a SAML message containing the previously generated authentication request.
- If the user is not already logged on to the IdP site or if re-authentication is required, the IdP asks for credentials (e.g., ID and password) and the user logs on.

6. Additional information about the user may be retrieved from the user data store for inclusion in the SAML response. (These attributes are predetermined as part of the federation agreement between the IdP and the SP—see [About attributes](#) on page 70 in the “Key Concepts” chapter of the PingFederate *Administrator's manual*.)
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 that POST responses be digitally signed.

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 scenario, the SP sends an authentication request to the IdP via HTTP POST. The returned SAML assertion is redirected through the user's browser. The response contains a SAML artifact .

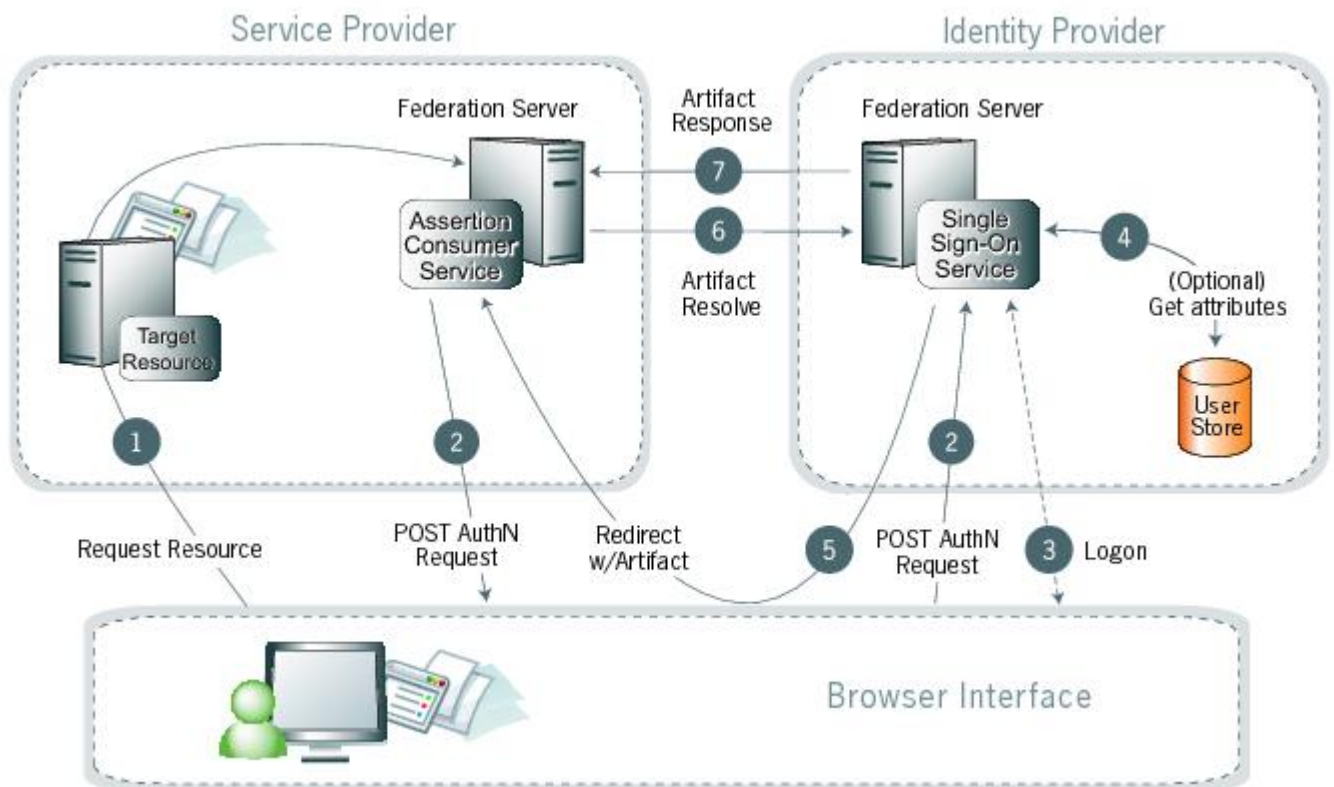


Figure 8: SP-Initiated SSO: POST/Artifact

Processing steps

1. A user requests access to a protected SP resource. The user is not logged on to the site. The request is redirected 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 is automatically posted to the IdP's SSO service.
3. If the user is not already logged on to the IdP site or if re-authentication is required, the IdP asks for credentials (e.g., ID and password) and the user logs on.
4. Additional information about the user may be retrieved from the user data store for inclusion in the SAML response. (These attributes are predetermined as part of the federation agreement between the IdP and the SP—see [About attributes](#) on page 70 in the “Key Concepts” chapter of the PingFederate *Administrator's manual*.)
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 a valid assertion is received, 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 SP sends an HTTP redirect message to the IdP containing a request for authentication. The IdP returns an artifact via HTTP redirect. The SP uses the artifact to obtain the SAML response.

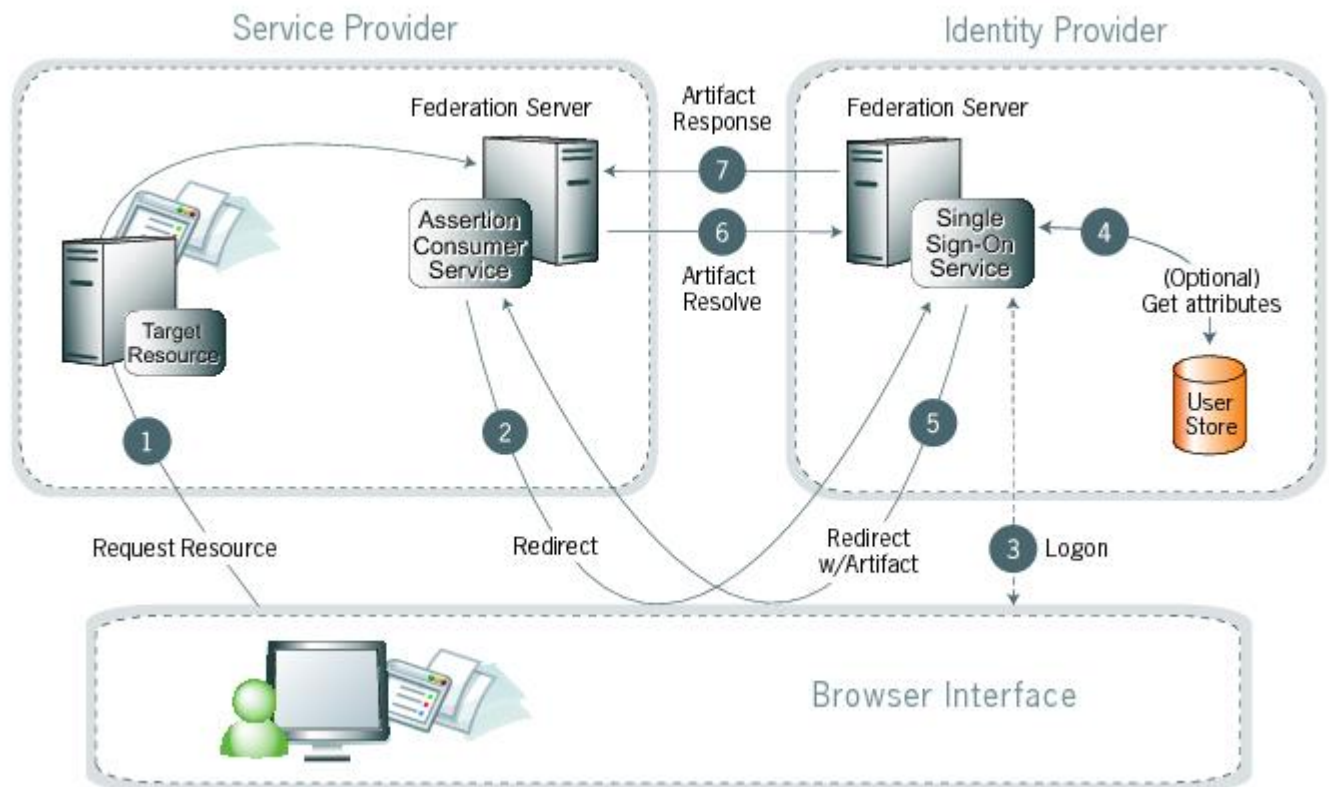


Figure 9: SP-Initiated SSO: Redirect/Artifact

Processing steps

1. A user requests access to a protected SP resource. The user is not logged on to the site. The request is redirected 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 SSO service.
3. If the user is not already logged on to the IdP site or if re-authentication is required, the IdP asks for credentials (e.g., ID and password) and the user logs on.
4. Additional information about the user may be retrieved from the user data store for inclusion in the SAML response. (These attributes are predetermined as part of the federation agreement between the IdP and the SP—see [About attributes](#) on page 70 in the “Key Concepts” chapter of the PingFederate *Administrator's manual*.)
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 ACS.
6. The ACS extracts the Source ID from the SAML artifact and sends an artifact-resolve message to the identity federation server's ARS.
7. The ARS sends a SAML artifact response message containing the previously generated assertion.

8. (Not shown) If a valid assertion is received, the SP establishes a session and redirects the browser to the target resource.

SP-initiated SSO—Artifact-Artifact

In this scenario, the SP sends a SAML to the IdP via 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.

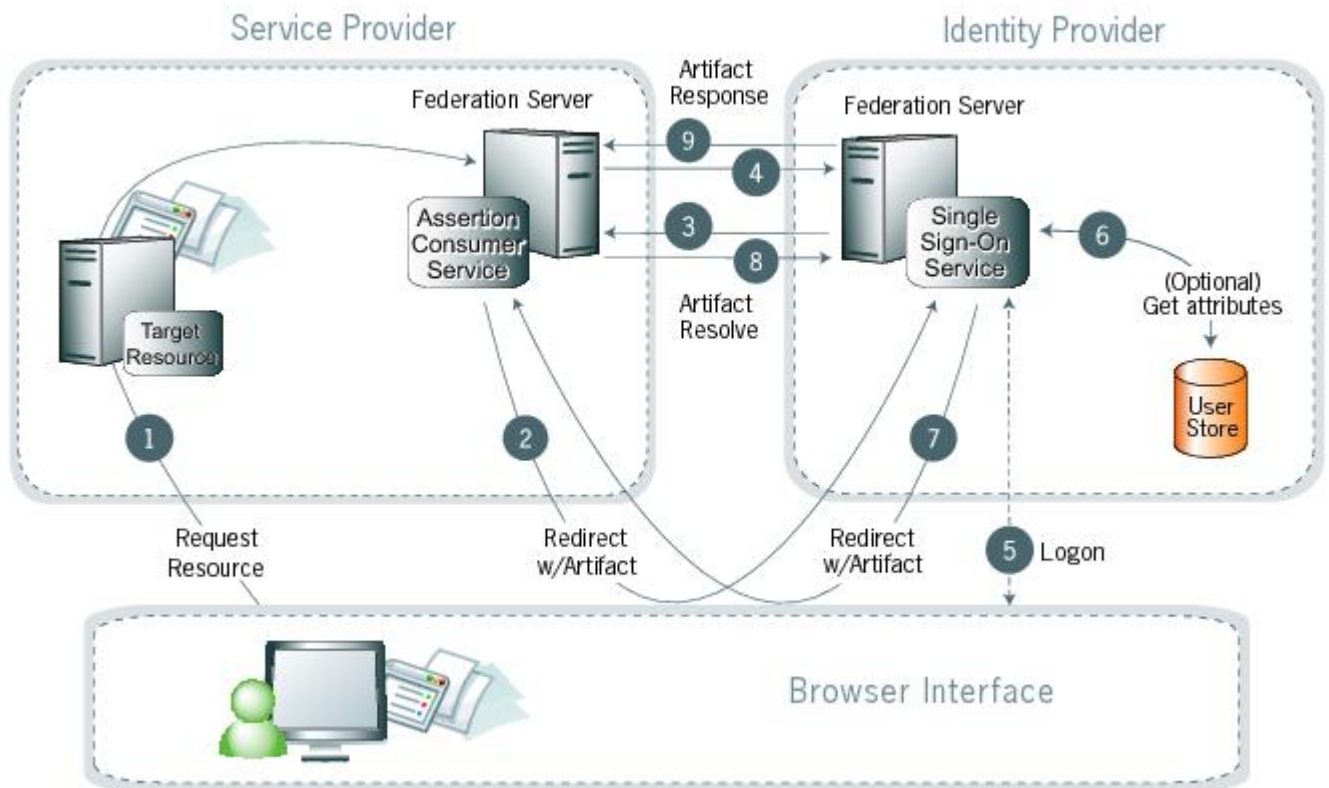


Figure 10: SP-Initiated SSO: Artifact/Artifact

Processing steps

1. A user requests access to a protected SP resource. The user is not logged on to the site. The request is redirected to the federation server to handle authentication.
2. The 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 SSO service.
 - Note:** The artifact contains the source ID of the SP's artifact resolution service 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 SP and IdP's source IDs and remote artifact resolution services are mapped according to the federation agreement 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 if re-authentication is required, the IdP asks for credentials (e.g., ID and password) and the user logs on.

6. Additional information about the user may be retrieved from the user data store for inclusion in the SAML response. (These attributes are predetermined as part of the federation agreement between the IdP and the SP—see [About attributes](#) on page 70 in the “Key Concepts” chapter of the PingFederate *Administrator's manual*.)
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 a valid assertion is received, the SP establishes a session and redirects the browser to the target resource.

IdP-initiated SSO—POST

In this scenario, 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 via HTTP POST.

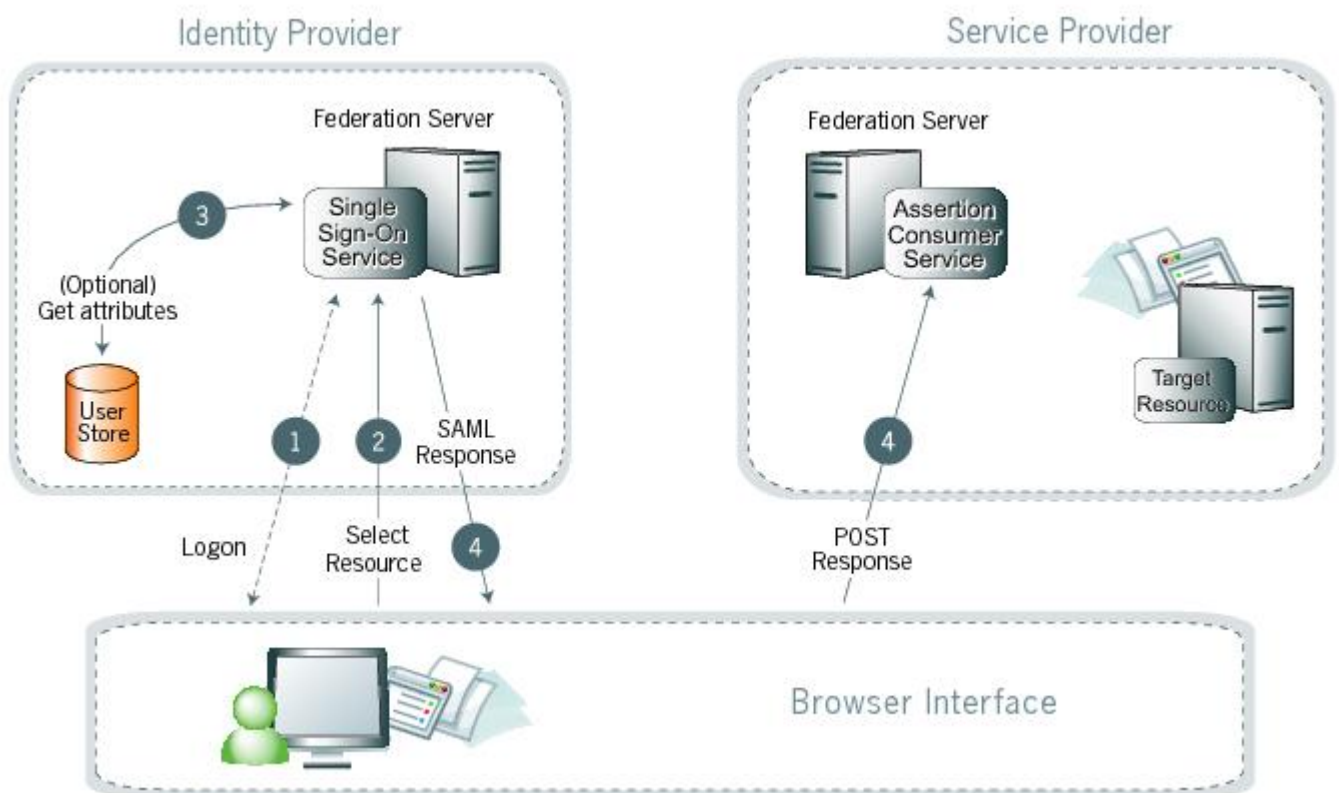



Figure 11: IdP-Initiated SSO: POST

Processing steps

1. A user has logged 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. Optionally, the IdP retrieves attributes from the user data store.
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 that POST responses be digitally signed.

- (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 scenario, the IdP sends a SAML artifact to the SP via an HTTP redirect. The SP uses the artifact to obtain the associated SAML response from the IdP.

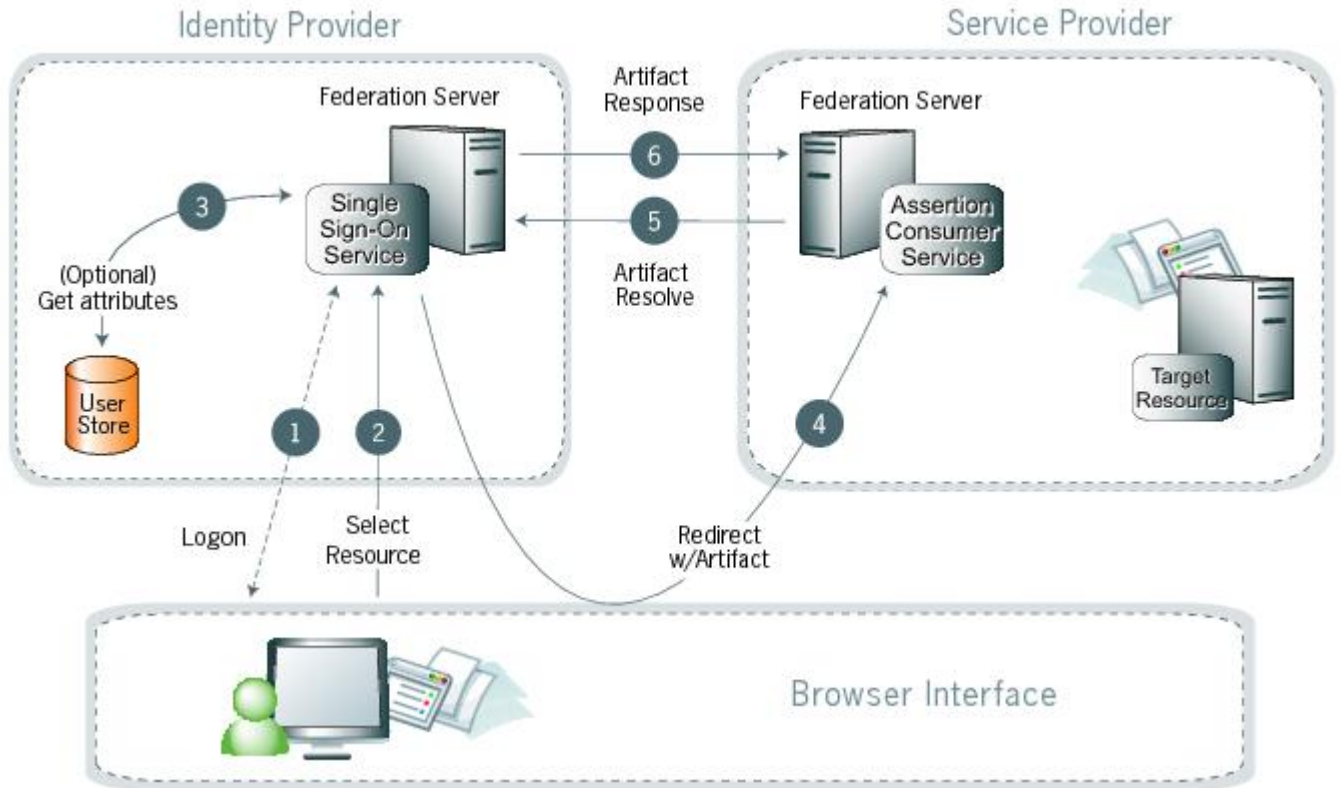


Figure 12: IdP-Initiated SSO: Artifact

Processing steps

- A user has logged 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).
- The user clicks a link or otherwise requests access to a protected SP resource.
- Optionally, the IdP retrieves attributes from the user data store.
- 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).
- 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).
- The ARS sends a SAML artifact response message containing the previously generated assertion.
- (Not shown) If a valid assertion is received, the SP establishes a session and redirects the browser to the target resource.

Single logout

The single logout (SLO) profile enables a user to log out of all participating sites in a federated session nearly simultaneously. The user may log out globally from any site, whether SP or IdP, as determined by respective web applications. The associated IdP federation deployment handles all logout requests and responses for participating sites. If a participating site returns an error, other participating sites may not receive their logout requests. In this scenario, PingFederate® returns an error message to the end users.


The logout messages may be transported using any combination of bindings described for SSO (POST, artifact, or redirect). Refer to the diagrams under [Single sign-on](#) on page 35 for illustrations of these message flows.

About session cleanup

When an SP receives an SLO request from an IdP, the session creation adapters must handle any session clean-up with respect to the local application.

Attribute Query and XASP

The SAML 2.0 Attribute Query profile allows an SP to request user attributes from an IdP in a secure transaction separate from SSO. The IdP, acting as an *Attribute Authority*, accepts Attribute Queries, performs a data-store 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.

 **Tip:** When privacy is required for sensitive attributes, you can configure PingFederate® to obfuscate (mask) their values in the server and transaction logs.

Web SSO is distinct from the Attribute Query use case; therefore, you can configure PingFederate servers to implement either or both of these profiles without regard to the other.


The X.509 Attribute Sharing Profile (XASP) defines a specialized extension of the general Attribute Query profile. The XASP specification enables organizations with an investment in PKI (Public Key Infrastructure) to issue and receive Attribute Queries based on user-certificate authentication.

Under XASP a user authenticates directly with an SP application by providing his or her 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 may be used to determine the correct IdP to use as the source of the requested attributes. Finally, the SP generates an Attribute Query and transmits it to the IdP over the SOAP back channel.

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

Standard IdP Discovery

SAML 2.0 IdP Discovery provides a cookie-based look-up mechanism used to identify a user's IdP dynamically during an SP-initiated SSO event, when the IdP is not otherwise specified. This mechanism can be helpful, in particular, in cases where an SP might be a hub for several IdPs in an identity federation.

 **Tip:** In addition to supporting standard IdP Discovery, PingFederate® provides a cross-protocol, proprietary mechanism allowing an 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 [Configure IdP discovery using a persistent cookie](#) on page 174.

In the standard scenario, when a user requests access to a protected resource on the SP, common-domain 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
- Common domain cookie reader

Each of these roles is necessary to support IdP Discovery. The roles may 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.

WS-Federation

PingFederate® supports the WS-Federation Passive Requestor Profile for SP-initiated SSO, enabling interoperability with Microsoft's Active Directory Federation Service (ADFS). This profile provides 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 SLO.



Note: Unlike SAML, WS-Federation consolidates the endpoints for SLO and SSO. So when you set up a WS-Federation connection in PingFederate, both types of transactions are available to an SP web application that supports them both.

For more information about WS-Federation and the Passive Requestor Profile, see [Web Services Federation Languages](http://docs.oasis-open.org/wsfed/federation/v1.2/os/ws-federation-1.2-spec-os.html) (docs.oasis-open.org/wsfed/federation/v1.2/os/ws-federation-1.2-spec-os.html).

Passive Requestor profile

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

Figure 19 illustrates message processing for SSO using WS-Federation.

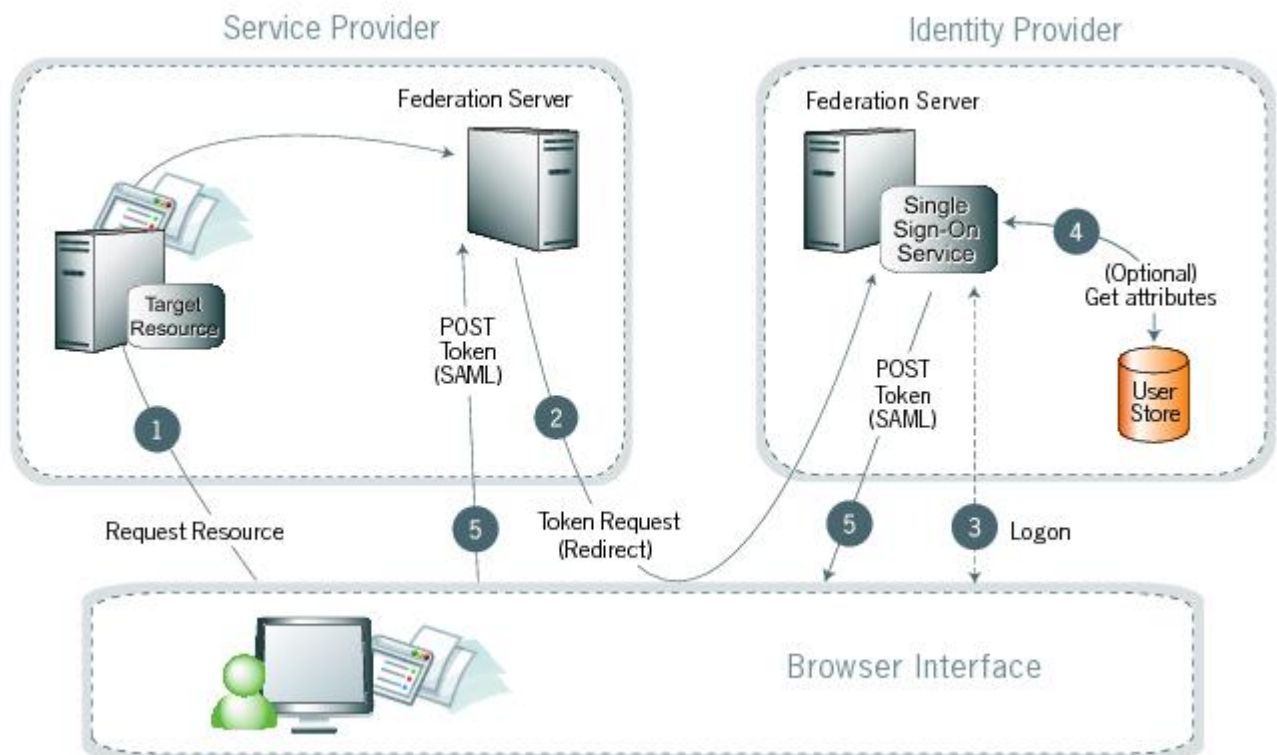


Figure 13: WS-Federation SSO

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 (e.g., ID and password) and the user logs on.
4. Additional information about the user may be retrieved from the user data store for inclusion in the SAML response. (These attributes are predetermined as part of the federation agreement between the IdP and the SP—see [About attributes](#) on page 70 in the “Key Concepts” chapter of the PingFederate *Administrator's manual*.)
5. The federation server creates a response containing a signed SAML assertion (or a JSON Web Token) and returns it to the SP via POST.
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.

Single logout using WS-Federation is handled in much the same way as with SAML (see [Single logout](#) on page 42); 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. This method of effectively interconnecting accounts across domains is supported by all protocols.

Account linking involves a *persistent name identifier* associated with accounts at each participating site. The name identifier, which may be an opaque pseudonym, is conveyed in the assertion. Once established locally, the SP can use the account link to look up the user and provide access without re-authentication.

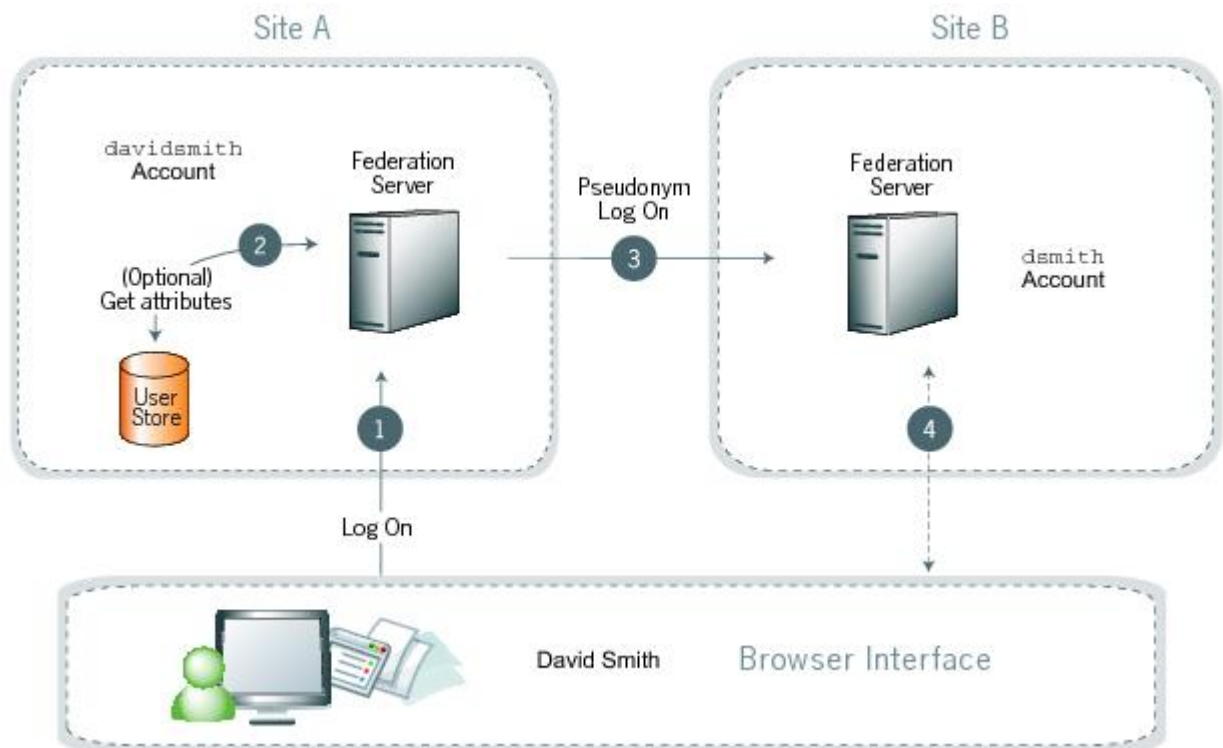


Figure 14: Account Linking

Processing steps

1. David Smith logs on to Site A as `davidsmith`. He then decides to access his account on Site B via Site A.
2. Optionally, the federation server looks up additional attributes from the data store.
3. The Site A federation server sends a persistent name identifier (possibly a pseudonym) to Site B, along with any other attributes.

If a pseudonym is used and other attributes are sent, care must be taken not to send attributes that could be used to identify the subject.

4. The federation server on Site B uses the information to associate the pseudonym with the existing account of dsmith. (Optionally, David is asked to provide consent to the linking.)

Once the link has been established, it is stored so that David only has to log on to Site A to have access to Site B.

Web Services standards

The PingFederate® WS-Trust STS is designed to interoperate with many different Web Service environments that support varying standards. PingFederate supports multiple versions of SOAP and WS-Trust specifications, and can freely operate with any combinations 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.)

Supported SOAP/WS-Trust versions and corresponding namespaces are listed in following table:

Table 1: 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.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, also WSSE) is a set of specifications defined by the Web Services Security Technical Committee (see www.oasis-open.org/committees/tc_home.php?wg_abbrev=wss) at the OASIS standards organization. WSS defines the XML extensions that can be used to secure Web Service invocations, providing a standard way for partners to add message integrity and confidentiality to their Web Service interactions. The WSS-defined token profiles describe standard ways of binding security tokens to these messages, enabling a variety of additional capabilities. The WSS technical committee has defined profiles for using SAML assertions, Username, Kerberos, X.509, and other existing security tokens. SSL/TLS is often used in conjunction with deployments of 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 46).



Figure 15: WSS Token Transfer

WS-Trust

WS-Trust comprises a protocol for systems and applications to use when requesting a service to issue, validate, and exchange security tokens. Organizations can leverage this protocol to centralize their security-token processing.

The WS-Trust specification also defines the role of a Security Token Service as the entity responsible for responding to requests using the protocol. In this role, the STS creates new security tokens, validates existing security tokens, and/or exchanges security tokens of one type for those of another (see [Figure 16: Token Exchange \(Example\)](#) on page 47).

WS-Trust was created by a consortium of leading platform and security vendors who have contributed the protocol to the OASIS standards organization, where it is managed by the WS-SX (Secure Exchange) technical committee (see www.oasis-open.org/committees/tc_home.php?wg_abbrev=ws-sx.)

Request types

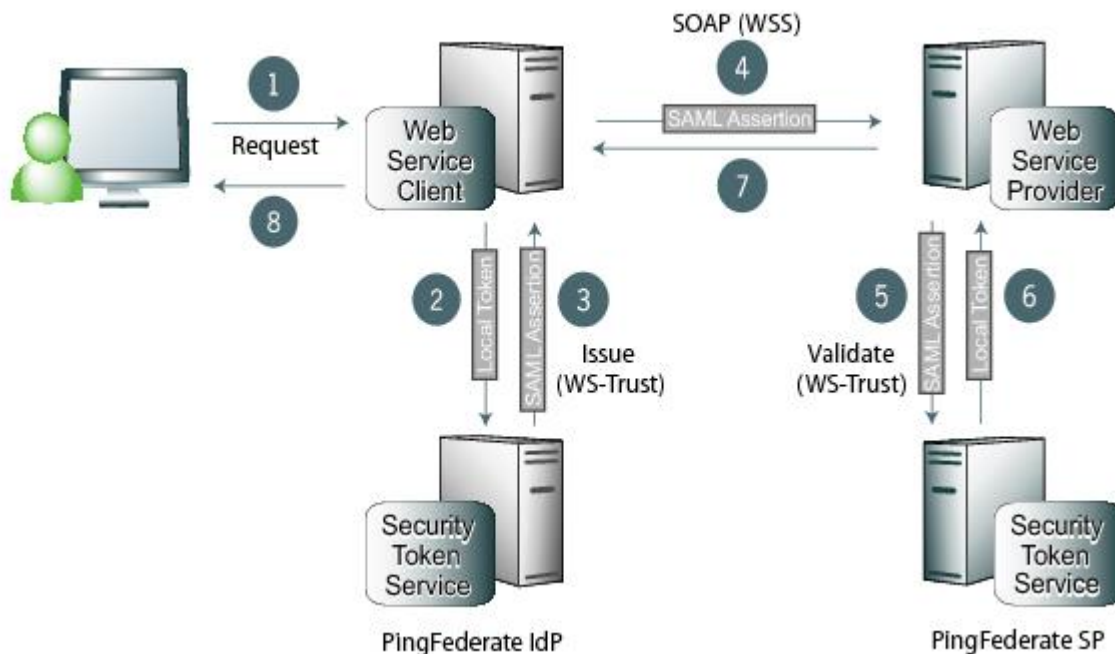
The WS-Trust protocol defines two request types that are particularly useful in securing Web Services: *Issue* and *Validate*, often associated with the Web Service Client (WSC) and Web Service Provider (WSP), respectively.

- The WSC requests that an 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 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).

The following illustration shows an example of a token exchange, using PingFederate to obtain a SAML assertion to be used in the WSS-secured Web Service call.

Figure 16: Token Exchange (Example)



Processing steps

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 WSS header.
5. The WSP retrieves the SAML assertion from the WSS 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 Representational State Transfer (REST) APIs (and non-REST APIs). Rather than 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 and/or policies appropriate to the client and the user. These tokens present less of a security and privacy risk than using secrets (or 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 properly 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. RO is the end user.

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. It 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](https://tools.ietf.org/html/rfc6749) (tools.ietf.org/html/rfc6749), the OAuth AS in PingFederate supports a wide variety of different interaction models appropriate for different types of clients such as a server, a desktop application, or an application on a phone or a tablet.

Web redirect flow

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 via an HTTP redirect. The client trades the authorization code for an access token, and then uses the token in a API call to obtain data.



Processing steps

1. User navigates to an OAuth client website (the requesting site) and requests access to protected resources from another website.

The OAuth client can optionally include the parameter `code_challenge` (with or without `code_challenge_method`) to reduce the risk of code interception attack. For more information, see step 3 and [Proof Key for Code Exchange by \(PKCE\) OAuth Public Clients](https://tools.ietf.org/html/rfc7636) (tools.ietf.org/html/rfc7636).

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

If the user is not already logged on, the OAuth AS challenges the user to authenticate. The OAuth AS authenticates the user and provides a **Request for Approval** page for the user to authorize the sharing of information.

Once the user authorizes, the OAuth AS redirects the browser to the requesting site with an authorization code.

If the user does not authenticate, an error is returned rather than the authorization code.

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

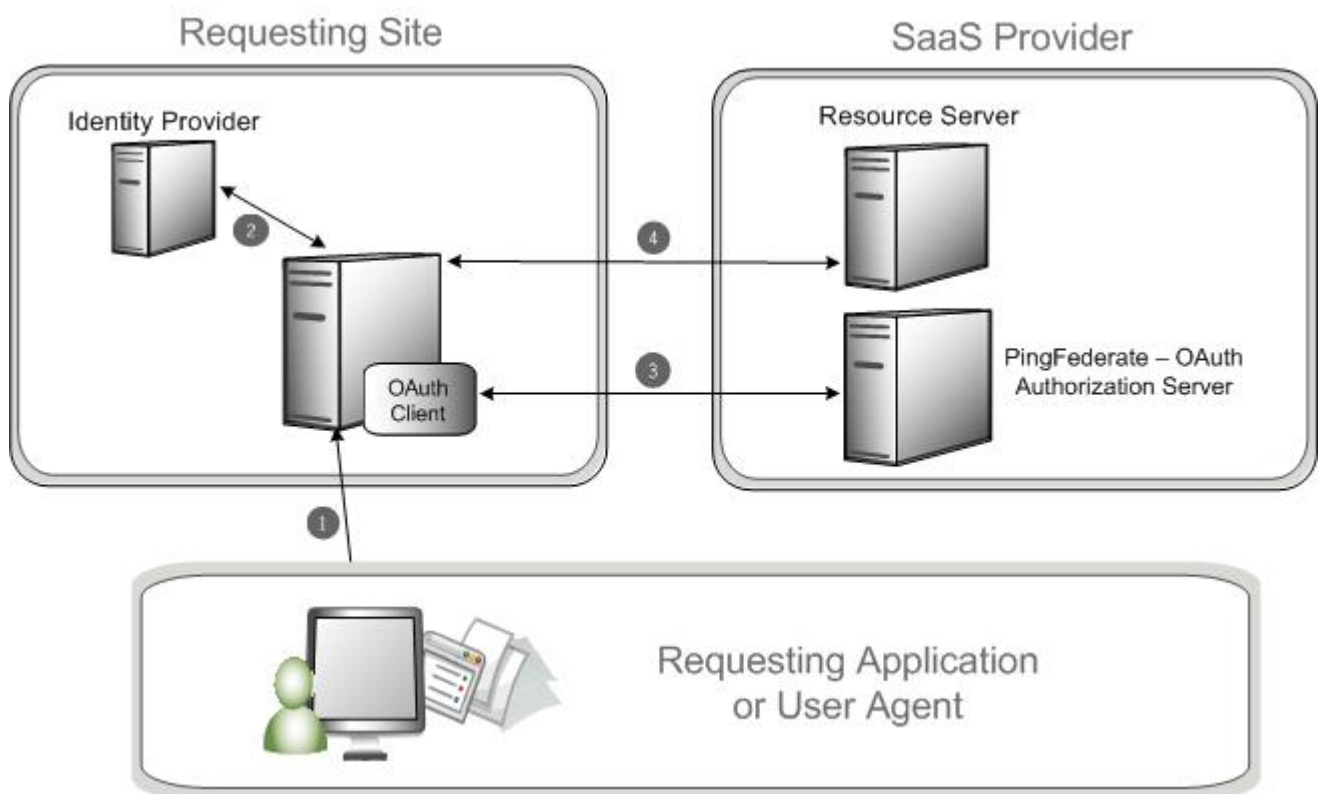
If the OAuth client has provided the optional parameter `code_challenge` in step 1, it must submit the corresponding `code_verifier` in this request.

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

4. The requesting site uses the access token in an API call to request user data.
5. The Resource Server 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 Resource Server returns the requested data to the requesting site.
7. (Not shown.) The requesting site displays data from the API call to the user.

SAML 2.0 profile for OAuth 2.0 authorization grants

In this scenario, a client obtains a SAML 2.0 bearer assertion and makes an HTTP request to the PingFederate® OAuth AS to exchange the SAML 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 Resource Server to obtain data.



Processing steps

1. Some user-initiated or client-initiated event (for example, 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 a SAML 2.0 bearer assertion from a local Identity Provider (IdP); for example, the PingFederate IdP server.
3. The client application makes an HTTP request to the PingFederate OAuth AS to exchange the SAML assertion for an access token. The OAuth AS validates the assertion and returns the access token. For more information on how the OAuth AS performs the validation, see the [specification](https://tools.ietf.org/html/rfc7522) (<https://tools.ietf.org/html/rfc7522>).
4. The client application adds the access token to its API call to the Resource Server. The Resource Server 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, a modern protocol for secure, lightweight transfer of authentication and user attributes (see openid.net/connect).

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

System for Cross-domain Identity Management (SCIM)

PingFederate® supports the SCIM 1.1 protocol for outbound and inbound provisioning. At an IdP (outbound) site, you can automatically provision and maintain user accounts at service-provider sites that have implemented SCIM. When PingFederate is configured as an SP (inbound), you can provision and manage user accounts and groups for your own

organization automatically using the standard SCIM protocol. The following table provides a brief summary of the supported features.

Feature	Outbound Provisioning	Inbound Provisioning
SCIM specification	SCIM 1.1	SCIM 1.1
Data format	JSON	JSON
User and group 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 client credentials	HTTP Basic and client certificate (mutual TLS)
Source data stores	Active Directory and other LDAPv3-compliant directory servers	Not applicable
Target data stores	Not applicable	Active Directory and custom data stores (via the Identity Store Provisioner Java SDK interface)

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

Transport and message security

The standards generally define two main ways of securing interactions: Secure Sockets Layer with Transport Level Security (SSL/TLS) and digital signatures. SSL/TLS is used in environments where both message confidentiality and integrity are required.

For SAML messaging, digital signatures are used to ensure the identity of both parties involved in the transaction and to validate that a message was received from a particular partner. With PingFederate®, you can also choose to encrypt SAML 2.0 messages, including SAML metadata files, as well as WS-Trust STS assertions to achieve increased privacy. For more information, refer to [Security and Privacy Considerations for the OASIS Security Assertion Markup Language \(SAML\) V2.0](https://docs.oasis-open.org/security/saml/v2.0/saml-sec-consider-2.0-os.pdf) (docs.oasis-open.org/security/saml/v2.0/saml-sec-consider-2.0-os.pdf).


Supported hardware security modules

For optimal security, PingFederate® can be configured 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.

PingFederate supports:

- SafeNet Luna SA HSM
- Thales nShield Connect HSM

Generally speaking, the first step is to install and configure the HSM according to the manufacturer's documentation. Once installed, follow the vendor-specific instructions to configure a new or an existing PingFederate to interact with the HSM for key generation, storage, and operation.

 **Tip:** Starting with PingFederate 8.3, you may enable the HSM hybrid mode, which provides you the choice to store each relevant key and certificate on the HSM or the local trust store. This capability allows your organization to transition the storage of keys and certificates to an HSM without the need to deploy a new PingFederate environment and to mirror the setup.

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. It is recommended that you consult your HSM vendor for performance tuning and optimization recommendations if you plan to use an HSM as part of your PingFederate deployment.

Install and configure Luna SA client and PingFederate

1. Install and configure your SafeNet Luna SA HSM, including the optional package for Java (referred to as the JSP), according to SafeNet's instructions.

This includes the creation of a partition, creation of a Network Trust Link (NTL), and assignment of a client to a partition. Ensure that you can perform the `vtl verify` command indicating that you are communicating securely and properly to the HSM.

Delete any unnecessary keys or objects that may have been created while testing communication to the HSM from the host that runs PingFederate®.

Note the password used to open communication to the HSM via the NTL. You need this for your PingFederate installation.

2. To enable the Java interface, copy the Luna library and program files to the Java installation as follows:

For 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.

For UNIX/Linux:

- Copy the `libLunaAPI.so` and `LunaProvider.jar` files from the `LUNA_HOME/jsp/lib` directory to the `JAVA_HOME/jre/lib/ext` directory.

SafeNet provides some sample Java applications that may be run to ensure that the Java HSM interface is working properly prior to installing PingFederate. Please contact SafeNet documentation for more information.

3. Edit the `JAVA_HOME/jre/lib/security/java.security` file in your Java environment and add the `LunaProvider` line to the list of security providers, *immediately before* the `sun.security.ec.SunEC` provider; for 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=com.safenetinc.luna.provider.LunaProvider
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. 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.

5. Edit the `hivemodule.xml` file in the `<pf_install>/pingfederate/server/default/conf/META-INF` directory and update the `<!-- Crypto provider -->` section.

a) For the `JCEManager` service endpoint, change the value of the construct class to as follows:

```
<construct class="com.pingidentity.crypto.LunaJCEManager5"/>
```

b) For the `CertificateService` service endpoint, change the value of the construct class as follows:

```
<construct class="com.pingidentity.crypto.LunaCertificateServiceImpl5"/>
```

6. Edit the `<pf_install>/pingfederate/bin/run.properties` file.

a) Change the value of the `pf.hsm.mode` property from `OFF` to `LUNA`.

b) If you are setting up a new PingFederate installation, set the value of the `pf.hsm.hybrid` property to `false`. When set to `false`, as you create or import certificates (such as your signing certificate or your encryption key), the certificates are stored on your HSM.

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 [Transition to an HSM](#) on page 193.

7. From the `<pf_install>/pingfederate/bin` directory, run the `hsmypass.bat` batch file for Windows or the `hsmypass.sh` script for UNIX/Linux.

Enter the NTL password when prompted (see [step 1](#)).

This procedure sets and securely stores the password for NTL communication to the HSM from PingFederate.



Note: The Luna SA HSM may be configured in a high-availability group. To do so, please refer to the SafeNet distributed-installation instructions. To properly synchronize data, ensure that the `HAOnly` property is enabled using this command: `vtl haAdmin -HAOnly -enable`

8. If you are setting up a new or configuring an existing PingFederate cluster, repeat these steps on each node.

This completes the steps required to configure PingFederate for use with the SafeNet Luna SA. You may start the new PingFederate server or restart the existing PingFederate server.



Important: To ensure expected behavior, SafeNet recommends restarting dependent processes such as PingFederate (including all server nodes in a cluster) whenever the Luna HSM is restarted.

Luna SA operational notes

Some restrictions apply to PingFederate® operations when using an HSM:

- When PingFederate is integrated with SafeNet Luna SA on a platform with Oracle Java SE Runtime Environment (Server JRE) 8 update 102, runtime errors may occur when handling certificates with a signing algorithm of RSA SHA256, SHA384, or SHA512. After switching to Oracle Server JRE 8 update 91, the issue is no longer observed.
- PingFederate does not store public certificates (for the purposes of signature verification, encryption, and back-channel authentication) on the hardware module. In this case, certificates are stored in the local trust store located on the file system.
- As an OpenID Provider, PingFederate generates and rotates temporary asymmetric key pairs to sign ID Tokens for Relying Parties. These in-memory short-term keys are not stored on the HSM.
- Private keys are not exportable. When configured for use with the HSM, administrative-console options for this feature are disabled. Only the public portion of generated keys is exportable.
- When running in FIPS 140-2 level 3 compliance (also known as strict FIPS mode) private keys can not be imported. In this case administrative-console options for this feature are disabled.
- Not all cipher suites in a standard Java configuration are available. They are limited to those listed in the `com.pingidentity.crypto.NcipherJCEManager.xml` file, located in the `<pf_install>/pingfederate/server/default/data/config-store` directory.
- 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, any deletion or creation of objects on the HSM not executed via the PingFederate user interface will not be recognized or operational.

For example, during the course of normal PingFederate operation 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 then delete object C and attempt to recover it via the data archive, PingFederate fails, producing various exceptions. Because the data archive contains a reference to the object and the object has been deleted from the HSM, it is not possible to use that data archive again.

Install and configure nShield Connect client and PingFederate

1. Install and configure your nShield Connect HSM client software.

As part of the installation, install the optional Java Support (including KeySafe) and nCipherKM JCA/JCE provider classes components.

2. After your installation, refer to the Thales nShield documentation to see how to make your PingFederate® server a client of an HSM server.



Note: PingFederate only supports Operator Card Set (OCS) protected keys. Note the password used for the OCS; you will need the password for your installation of PingFederate.

3. If you have not already done so, download and install the (JCE) Unlimited Strength Jurisdiction Policy Files 8 (jce_policy-8.zip).

Follow instructions in the readme to install the Policy Files.

4. To enable the Java interface, copy the `NFAST_HOME/java/classes/nCipherKM.jar` file to the `JAVA_HOME/jre/lib/ext` directory.

Thales provides some sample Java applications that may be run to ensure that the Java HSM interface is working properly prior to installing PingFederate. Please refer to Thales documentation for more information.

5. Edit 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* all the `sun` providers; for example:

```
# List of providers and their preference orders (see above):
security.provider.1=sun.security.provider.Sun
security.provider.3=com.sun.net.ssl.internal.ssl.Provider
security.provider.4=com.sun.crypto.provider.SunJCE
security.provider.5=sun.security.jgss.SunProvider
security.provider.6=com.sun.security.sasl.Provider
security.provider.7=org.jcp.xml.dsig.internal.dom.XMLDSigRI
security.provider.8=sun.security.smartcardio.SunPCSC
security.provider.9=sun.security.mscapi.SunMSCAPI
security.provider.10=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. Edit the `hivemodule.xml` file in the `<pf_install>/pingfederate/server/default/conf/META-INF` directory and update the `<!-- Crypto provider -->` section.

- a) For the `JCEManager` service endpoint, change the value of the `construct class` to as follows:

```
<construct class="com.pingidentity.crypto.NcipherJCEManager"/>
```

- b) For the `CertificateService` service endpoint, change the value of the `construct class` as follows:

```
<construct class="com.pingidentity.crypto.NcipherCertificateServiceImpl"/>
```

8. Edit the `<pf_install>/pingfederate/bin/run.properties` file.

- a) Change the value of the `pf.hsm.mode` property from `OFF` to `NCIPHER`.

- b) If you are setting up a new PingFederate installation, set the value of the `pf.hsm.hybrid` property to `false`. When set to `false`, as you create or import certificates (such as your signing certificate or your encryption key), the certificates are stored on your HSM.

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 [Transition to an HSM](#) on page 193.

9. From the `<pf_install>/pingfederate/bin` directory, run the `hsmpass.bat` batch file for Windows or the `hsmpass.sh` script for UNIX/Linux.

Enter the Operator Card Set password when prompted (see [step 2](#)).

This procedure sets and 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, locate 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 perform define this environment variable on all servers within the cluster.

- d) Restart the nShield Connect hardserver on all PingFederate servers in the cluster. (See the Thales documentation for instructions on restarting the hardserver.)
- e) Log on to the PingFederate administrative console, go to the **Server Configuration > Cluster Management** screen and then click **Replicate Configuration** to push the configuration changes, which includes the nShield data, to the engine nodes.

This completes the steps required to configure PingFederate for use with the Thales nShield Connect. You may start the new PingFederate server or restart the existing PingFederate server.

- ❗ **Important:** To ensure expected behavior, PingFederate (including all server nodes in a cluster) should be restarted whenever the nShield HSM is restarted.

nShield Connect operational notes

Some restrictions apply to PingFederate® operations when using an HSM:

- When PingFederate is integrated with Thales nShield Connect on a platform with Oracle Java SE Runtime Environment (Server JRE) 8 update 102, runtime errors may occur when handling certificates with a signing algorithm of RSA SHA256, SHA384, or SHA512. Upgrading to Oracle Server JRE 8 update 112 resolves these runtime errors.
- PingFederate only supports Operator Card Set (OCS) protected keys. If you use a standard (non-persistent) OCS, the HSM removes the protected keys from its memory when the card is removed from the smart card reader. Requests will likely fail because almost all requests require cryptographic processing. To resume operations, you must insert the card back to the smart card reader and then restart PingFederate.

Alternatively, you may 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. Note that no new keys and certificates can be created and stored on the HSM until the card is inserted back to the HSM. (No restart of PingFederate is required.) For more information about persistent OCS, please consult your HSM vendor.

- As an OpenID Provider, PingFederate generates and rotates temporary asymmetric key pairs to sign ID Tokens for Relying Parties. These in-memory short-term keys are not stored on the HSM.

- Private keys are not exportable. When configured for use with the HSM, administrative-console options for this feature are disabled. Only the public portion of generated keys is exportable.
- When running in FIPS 140-2 level 3 compliance (also known as strict FIPS mode) private keys can not be imported. In this case administrative-console options for this feature are disabled.
- Not all cipher suites in a standard Java configuration are available. They are limited to those listed in the `com.pingidentity.crypto.NcipherJCEManager.xml` file, located in the `<pf_install>/pingfederate/server/default/data/config-store` directory.
- 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, any deletion or creation of objects on the HSM not executed via the PingFederate user interface will not be recognized or operational.

For example, during the course of normal PingFederate operation 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 then delete object C and attempt to recover it via the data archive, PingFederate fails, producing various exceptions. Because the data archive contains a reference to the object and the object has been deleted from the HSM, it is not possible to use that data archive again.

Administrator's manual

This manual provides information about using Ping Identity®'s PingFederate® to deploy a secure Internet single sign-on (SSO) solution based on the latest security and e-business standards.

The Administrator's manual consists of:

- [Key concepts](#) on page 57 — A discussion of central concepts needed to understand Internet SSO, the WS-Trust Security Token Service (STS), and PingFederate deployment and administration.
- [System administration](#) on page 90 — Information about maintaining the PingFederate server and deployment, using log files, managing users, and handling other administrative functions.
- [System settings](#) on page 145 — How to configure your local PingFederate server settings.
- [OAuth configuration](#) on page 241 — How to configure PingFederate to act as an OAuth Authorization Server.
- [Security management](#) on page 178— Information about importing, exporting, and maintaining certificates and keys in PingFederate.
- [Authentication policies](#) on page 209— Configure Authentication Policies, Selectors, and Policy Contracts to fulfill various authentication requirements, to deploy PingFederate as a Federation Hub, or both.
- [Identity provider SSO configuration](#) on page 275— How to configure PingFederate to act as an Identity Provider (IdP) and establish connections to Service Providers.
- [Service provider SSO configuration](#) on page 330 — How to configure PingFederate to act as a Service Provider (SP) and establish connections to Identity Providers.
- [WS-Trust STS configuration](#) on page 400 — How to configure PingFederate to act as a Security Token Service for Web Service Clients and Providers in either an IdP or SP environment.
- [IdP-to-SP bridging](#) on page 419 — Configure advanced federation use cases, such as Adapter-to-Adapter Mappings (for Browser SSO) and Token Exchange Mappings (for STS).
- [Kerberos Adapter](#) on page 427 — How to configure PingFederate to use the packaged Kerberos Adapter.
- [HTML Form Adapter](#) on page 430 — How to configure PingFederate to use the packaged HTML Form Adapter.
- [HTTP Basic Adapter](#) on page 438 — How to configure PingFederate to use the packaged HTTP Basic Adapter.
- [OpenToken Adapter](#) on page 440 — How to configure PingFederate to use the packaged OpenToken Adapter for interfacing with your web applications.
- [Composite Adapter](#) on page 445 — How to configure PingFederate to use the packaged Composite Adapter.
- [Application endpoints](#) on page 448 — Detailed information about using PingFederate connection endpoints for web single sign-on and single logout.
- [OAuth 2.0 endpoints](#) on page 463 — Detailed information about using PingFederate connection endpoints for the OAuth.
- [Web Service interfaces](#) on page 477 — Information developers can use to automate connection-configuration management and runtime SSO partner discovery.
- [Attribute mapping expressions](#) on page 497 — How to enable and use expressions in conjunction with mapping attributes.
- [Customize assertions and authentication requests](#) on page 501 — Information about assertion and authentication request customizations.
- [Fulfillment by data store queries](#) on page 506 — How to configure data store queries to fulfill a contract or to pull attributes to be verified in the token authorization process
- [Troubleshooting](#) on page 513 — Solutions for difficulties that may be encountered.
- [Glossary](#) on page 522 — Definitions of terms used in the manual and in identity federation parlance.
- [List of acronyms](#) on page 531 — A list of acronyms.

Key concepts

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

- [Connection types](#) on page 58
- [About WS-Trust STS](#) on page 58
- [About OAuth](#) on page 61
- [SSO integration kits and adapters](#) on page 67
- [Hierarchical plug-in configurations](#) on page 69
- [Identity mapping](#) on page 69
- [About attributes](#) on page 70
- [Security infrastructure](#) on page 75
- [Using Auto-Connect](#) on page 79
- [User provisioning](#) on page 81
- [Federation planning checklist](#) on page 82

 **Tip:** For an introduction to secure single sign-on (SSO), federated identity management, and Ping Federate product features, see [About identity federation and SSO](#) on page 9 in the *Get started with PingFederate* guide.

Connection types

PingFederate features an integrated administrative console for configuring four kinds of connections to identity-federation partners:

- Browser-based SSO – Also called Browser SSO in the administrative console, this term is often used to refer to standards-based secure SSO, which generally depends on a user's browser to transport identity assertions and other messaging between partner endpoints (see [Supported standards](#) on page 29 in the *Get started with PingFederate* guide).
- WS-Trust STS – Employs the PingFederate Security Token Service (STS), which enables Web Service Client and Provider applications to extend SSO to identity-enabled Web Services at provider sites, using another set of standards (see the next section, [About WS-Trust STS](#) on page 58). 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) (see [About OAuth](#) on page 61).
- Provisioning – Provides automated cross-domain inbound and outbound user management (see [User provisioning](#) on page 81).

The types of connections can be configured together for the same partner or independently.

About WS-Trust STS

The PingFederate WS-Trust STS allows organizations to extend SSO identity management to Web Services. (For information about WS-Trust and the role of an STS, see [Web Services standards](#) on page 46 in the “Supported Standards” chapter of the *Get started with PingFederate* guide.)

The WS-Trust STS can be configured for partner connections independently or in conjunction with browser-based SSO for either an IdP or an SP deployment. The STS is bundled with separate plug-ins for standard SAML (Security Assertion Markup Language) token processing and generation (see [Token processors and generators](#) on page 59).

Connection-based policy

For both the IdP and SP roles, PingFederate employs a partner-connection configuration, which enables the association of Web Services authentication policies with federation partners. For STS processing, these policies define configurations for handling WS-Trust requests and transferring identity information between security domains (see [Web Services standards](#) on page 46 in the “Supported Standards” chapter of the *Get started with PingFederate* guide).

IdP configuration

In an IdP role, you use the administrative console to configure WS-Trust request-processing policy for your SP partner including:

- The type of SAML token to create—suitable for consumption by the intended Web Service Provider (WSP, at the SP site)—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

In an SP role, you use the administrative console 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 that plug into the PingFederate server. These plug-ins deploy similarly to browser-based SSO adapters (see [SSO integration kits and adapters](#) on page 67).

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

For an 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.

Only SAML 1.1 or 2.0 tokens are generated by PingFederate configured as an IdP for sending across trust boundaries to a federated SP partner. Likewise, only SAML tokens are accepted by PingFederate configured as an SP. Token plug-ins allow a modular approach for validating and producing the various token types used by different applications or systems within a conceptual trust domain. PingFederate provides bundled and separately available token plug-ins.

 **Tip:** For direct STS token exchange within the same domain or trust boundary, you can also use the PingFederate STS to exchange one token type for another directly, without generating a transitional SAML token (see [Token translator mappings](#) on page 423).

PingFederate also allows you to use a configuration of a token processor or generator as a parent instance from which you can create child instances (see [Hierarchical plug-in configurations](#) on page 69).

Bundled token plug-ins

PingFederate is installed with token processors for an IdP configuration that accept and validate SAML 1.1, 2.0 tokens, OAuth Bearer access tokens, JWT tokens, Username tokens, and Kerberos tokens (see [Token models and management](#) on page 62). (SAML tokens are issued on the IdP side via built-in browser-based SSO capabilities.)

For an SP configuration, token generators are provided for issuing local SAML 1.1 or 2.0 tokens. (Incoming SAML tokens are validated, once again, by using built-in capabilities.)


Commercial token plug-ins

Ping Identity provides token plug-ins to work with various authentication systems and leading identity management systems. Available plug-ins, together known as *Token Translators*, may be downloaded from the Ping Identity [Downloads](#) website.

WSC and WSP support

Ping Identity provides the Java client Software Development Kit (SDK) for enabling Web Service applications (WS clients or providers) to interact with the PingFederate STS.

In addition, for 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 at 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 access-policy decisions, when specifically categorized user attributes are sent in the security token (see [STS namespaces](#) on page 71).

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 may be necessary, the PingFederate STS supports the use of customBinding.

For more information about bindings, see Microsoft's [System-Provided Bindings](#) (docs.microsoft.com/en-us/dotnet/framework/wcf/system-provided-bindings).

Developers can obtain metadata from PingFederate to expedite configuring their applications. PingFederate offers two varieties of metadata, which are often used together to arrive at functional WSC and WSP configurations:

- Federation Metadata, which contains details on the PingFederate public signing certificate and other information required to establish the trust relationship (see [/pf/federation_metadata.ping](#) on page 462).
- Metadata Exchange, which contains connection details relating to the SP partner (see [/pf/sts_mex.ping](#) on page 462).

For more information about claim-based federated identity, see Microsoft's [A Guide to Claims-based Identity and Access Control](#) (msdn.microsoft.com/library/ff423674.aspx).

STS OAuth integration

PingFederate STS provides several ways to facilitate the use of issued tokens with an OAuth 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 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 (tools.ietf.org/html/rfc7523)).

OAuth Access Token via JWT Bearer Token grant type

```
oauth-v2:access:token:response|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 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 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 (tools.ietf.org/html/draft-ietf-oauth-saml2-bearer).

OAuth Access Token via 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 may have with an OAuth AS, allowing the client to obtain additional resources on behalf of already-authenticated users in follow-on transactions.

About OAuth

PingFederate® can be configured to act as an OAuth Authorization Server (AS), allowing a resource owner (typically, an end user) to grant authorization to an OAuth Client requesting access to resources hosted by a Resource Server (RS). The OAuth AS issues tokens to clients on behalf of a resource owner for use in authenticating a subsequent API call to the RS—typically, but not exclusively, a REST API call.

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

The PingFederate OAuth AS can be configured independently or in conjunction with STS or browser-based SSO for either an IdP or an SP deployment.

In an IdP deployment, an IdP adapter can be used to authenticate and provide user information for the access token.

In an SP deployment, the inbound SAML assertion can be used to provide authentication information about the user that can be associated with the access token through an OAuth attribute mapping in the IdP connection.

For an STS IdP, an OAuth token processor is provided with the PingFederate installation to validate incoming OAuth Bearer access tokens.

Delegated access types

Explicit delegation

This is the most common OAuth use case, which involves a resource owner (RO) who explicitly delegates to a client the authority 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 AS to issue access tokens for use in authenticating an API call. These tokens may be characterized by both their security model and data model.

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 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 (JavaScript Object Notation), signed by the AS and verified directly by the RS.

Reference tokens (Internally Managed Reference Tokens)

Serve as a reference to some set of attributes. The RS must de-reference the token for the corresponding identity and security information at the OAuth AS that issued it.

Token management

PingFederate[®] supports multiple access token management instances, providing flexibility for enterprises where deployments may 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 AS, sending a request for an access token that includes an access grant. An access grant is also used when an RS requests validation of an access token from the OAuth AS.

Primary grant types

OAuth defines several different access grant types—each reflecting different authorization mechanisms.

Authorization Code

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

Resource Owner Credentials

The client collects the resource owner's password and exchanges it at the OAuth 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

A refresh token is often returned with an access token. Once the original access token expires, the corresponding refresh token can be sent to the OAuth AS to obtain a fresh access token without requiring the resource owner to reauthenticate. 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 can only be used in conjunction with either the Authorization Code or Resource Owner Password Credentials grant type.

Implicit

An access token is returned to the client through a browser redirect in response to the resource owner authorization request (rather than an intermediate authorization code). This grant type is suitable for clients incapable of keeping client credentials confidential (for use in authenticating with the OAuth AS) such as client applications implemented in a browser using a scripting language like JavaScript.

Client Credentials

The client presents its own credentials to the OAuth AS in order to obtain an access token. This access token is either associated with the client's own resources, and not a particular resource owner, or is associated with a resource owner 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:

Validation Grant Type

```
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 the OAuth AS. The grant type allows an RS to check with the OAuth AS on the validity of a bearer access token received from a client making a protected-resources call.

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 OAuth AS. 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 STS processing.

Scopes

Where OAuth provides a mechanism to constrain the privileges associated with an access token, scopes provide a way to more specifically define the privileges requested and granted.

Generally, a client specifies the scopes desired when asking for authorization. The issued access token is associated with these approved scopes. Scopes are configured globally in the **OAuth Settings > Scope Management** screen but may be restricted on a client-by-client basis.

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 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 can authenticate via client secret and private key JWT. It can also handle the scenario where authentication is not required.

Persistent vs. transient grants

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

Transient grants are valid only for the lifetime of the access token itself. The **Client Credential** access grants, for example, are considered transient.

Persistent grants can result from the following OAuth use cases:

- OAuth clients using the **Refresh Token** grant type in conjunction with either the **Authorization Code** or **Resource Owner Credentials** grant type.

If the use cases involve mapping attributes from authentication sources (IdP adapter instances or IdP connections) or Password Credential Validator (PCV) instances to the access tokens (directly or through persistent grant extended attributes), such attributes and their values are stored along with the persistent grants so that they can be reused when clients subsequently present refresh tokens for new access tokens.


- OAuth clients using the **Implicit** grant type *and* the **Reuse Existing Persistent Access Grants for Grant Types** check box is selected in the **OAuth Settings > Authorization Server Settings** screen.

If the use cases involve mapping attributes from authentication sources or PCV instances to the access tokens (directly or through persistent grant extended attributes), attribute values are obtained at runtime for each token request. No attributes or their values are stored with the persistent grants.


Persistent grants (and the associated attributes and their values, if any) remain valid until the grants expired or are explicitly revoked.

Support for persistent grants requires PingFederate® to use a database server or an LDAP directory server for long-term storage. The data structure contains a `USER_KEY` attribute, which can be populated according to information mapped using attributes obtained during a user's initial authentication verification within PingFederate.

PingFederate automatically removes expired grants and the associated attributes once a day. As needed, you can fine-tune the frequency and the size of the cleanup batch.

 **Important:** PingFederate uses a pre-installed HyperSQL database as its grant data store for initial setup and testing. We however strongly recommend that you choose your own, secured database for production deployments.

PingFederate also supports other storage solutions through the PingFederate SDK. For more information, see the Javadoc for the `AccessGrantManager` interface.

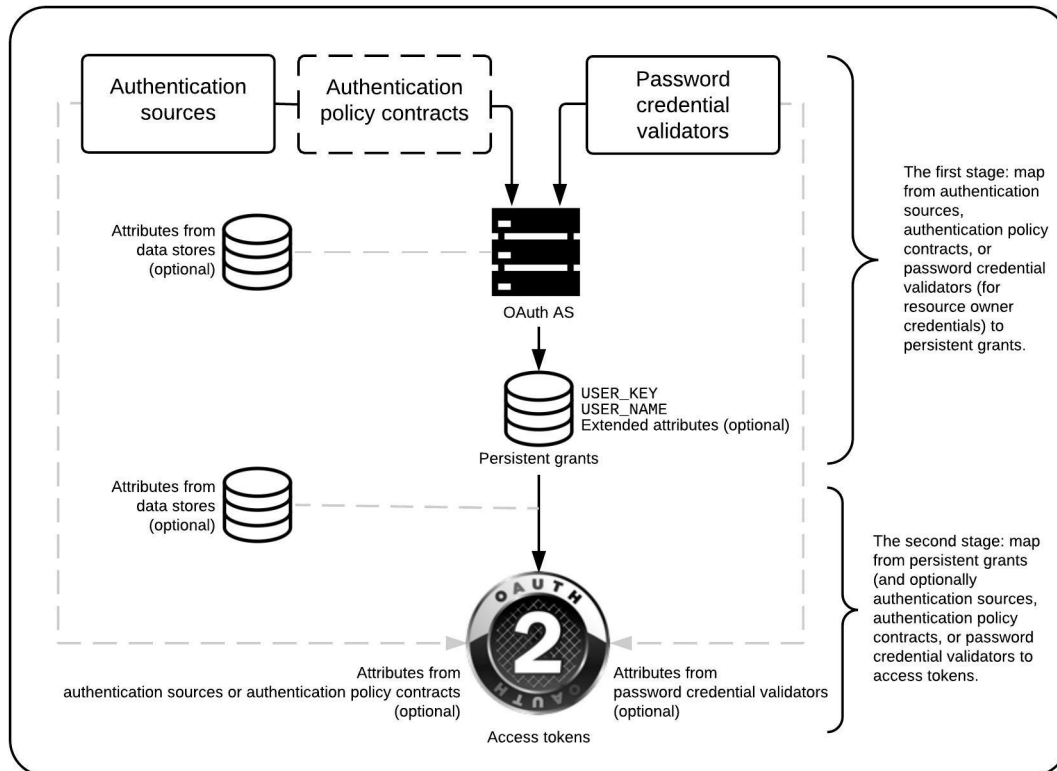
 **Tip:** The Javadoc for PingFederate is located the `<pf_install>/pingfederate/sdk/doc` directory.

Mapping OAuth attributes

Mapping OAuth attributes is a two-stage processing workflows:

- The first stage: map from authentication sources, authentication policy contracts, or password credential validators (for resource owner credentials) to persistent grants.
- The second stage: map from persistent grants (and optionally authentication sources, authentication policy contracts, or password credential validators) to access tokens.

Note that 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, mapping is required for setting up persistent grants, including a user key and all extended attributes.

The mappings may use attributes obtained during initial authentication events within PingFederate, namely attributes from IdP adapter instances, attributes from assertions via IdP connections, attributes from authentication policy contracts, or attributes returned by password credential validator instances. Moreover, data store queries may also be configured; for example, to retrieve the user identifier from an LDAP directory server as the user key.



Note: The USER_KEY values must be unique across all end users because it is the identifier to store and to retrieve persistent grants. For example, if you have two Active Directory domains, the sAMAccountName value of an end user in one domain may collide with that of another end user in the other domain. In this scenario, you can map Subject DN to the USER_KEY.

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.

You may also set up specific mappings between the authentication sources, authentication policy contracts, or password credential validators. When the authentication context matches a specific mapping, attributes from the authentication sources, authentication policy contracts, or password credential validators can be mapped into OAuth access tokens.


Data stores may also be used here to retrieve any required user attributes.

Runtime processing

At runtime, the first time a client requests an OAuth token the two mapping sequences are employed sequentially. The second mapping is invoked every time a new access token is requested based on an existing persistent grant.

OAuth user-facing screens

The PingFederate OAuth AS provides two screens presented to end users (resource owners) during OAuth transactions, one requesting approval of the scope of protected resources requested and one providing a means of revoking persistent access grants.

 **Tip:** These screens can be customized and branded as needed (see [Customizable user-facing screens](#) on page 126).

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 (see openid.net/connect).

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 is the release of at least two tokens to the requesting client application: an ID token and an OAuth access token. (Depending on associated grant types, a refresh token may also be released.)

The ID token is an integrity-secured, self-contained token in JSON Web Token (JWT) format containing claims about the user, namely the subject. A client uses the ID token to securely identify the user authenticated by an OP accessing the client application. 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`).

To configure OpenID Connect policies, first enable the protocol in **Server Configuration > Server Settings > Roles & Protocols** screen. To make use of OpenID Connect defined scopes for accessing various degrees of user-associated claims, you must also configure them in the **OAuth Settings > Scope Management** screen.

 **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.

Furthermore, 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 `/pf-ws/rest/sessionMgmt/revokedSris`).

As an OP, PingFederate can optionally accept request parameters via self-contained, signed JWTs. This capability enables PingFederate to validate the integrity of the request parameters it receives before processing the request further. For more information, see the section about passing a request object by value in the OpenID Connect specification at http://openid.net/specs/openid-connect-core-1_0.html#RequestObject.

Relying Party support

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

The setup involves establishing an IdP connection to the OP. In essence, PingFederate retrieves identity information from the OP and passes the end-user claims, which are basically 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, a modern standard that has been gaining momentum in the industry.

PingFederate is also capable of sending request parameters via self-contained, signed JWTs, thus adding a layer of security to the transmission of the request parameters.

SSO integration kits and adapters

As a stand-alone 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 SSO.



Note: See the PingFederate [SSO integration overview](#) on page 552 for more information.

For an IdP (the first mile), this integration process involves providing a mechanism through which PingFederate can look up a user's current authenticated session data (for example, a cookie) or authenticate a user without such a session. For an 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 commercial 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, for IdP and federation hub deployments PingFederate provides plug-in *authentication selectors*, which enable dynamic selection of authentication sources (IdP adapter instances, or IdP connections for federation hub use cases) based on administrator-specified criteria (see [Bundled authentication selectors](#) on page 67).

PingFederate also includes a robust software development kit (SDK), which software developers can use to write their own adapters for specific systems. Adapters can be written to retrieve attributes from custom data stores, connect to application- or IdM-specific user authentication systems, or provide complex attribute transformations or processing.

Bundled adapters

PingFederate comes bundled with a set of adapters.

Kerberos Adapter

Provides a seamless desktop SSO experience for Windows environments and supports authentication mechanism assurance from Active Directory domain service. This adapter is recommended for new configurations as a simpler alternative to the separately available IWA Integration Kit (see [Kerberos Adapter](#) on page 427).

HTML Form Adapter and HTTP Basic Adapter

Used in conjunction with Password Credential Validators (see [Manage password credential validators](#) on page 196). These adapters provide integration, for example, with LDAP user-data stores such as Active Directory or direct user logon via credentials maintained by PingFederate (see [HTML Form Adapter](#) on page 430 and [HTTP Basic Adapter](#) on page 438).

OpenToken Adapter

Provides a generic interface for integrating with various applications, including Java- and .NET-based applications (see [OpenToken Adapter](#) on page 440).

Composite Adapter

Allows multiple configured IdP adapters to execute in sequence. This capability, called *adapter chaining*, may be used either for single-adapter usage, depending on authentication context, or to support multifactor authentication via a series of adapters.

Bundled authentication selectors

For IdP and federation hub deployments, PingFederate also includes global (cross-connection) authentication selectors. Along with the Composite Adapter (see the previous section) and token authorization, the selectors enable dynamic integration with an organization's authentication or authorization policies (also known as *adaptive federation*).



Tip: The results of authentication-selection criteria evaluation may be used to select subsequent selectors or authentication sources, which allows handling of complex hierarchical access-policy decisions (see [Authentication policies](#) on page 209).

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 a specified range, or 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 (see [Configure the CIDR Authentication Selector](#) on page 210).

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, this selector allows you to choose whether or not Integrated Windows Authentication is attempted based on the PingFederate cluster node with which a Key Distribution Center is associated (see [Configure the Cluster Node Authentication Selector](#) on page 211).

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 (see [Configure the Connection Set Authentication Selector](#) on page 211).

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 (see [Configure the HTTP Header Authentication Selector](#) on page 212).

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 (see [Configure the HTTP Request Parameter Authentication Selector](#) on page 213).

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 (see [Configure the OAuth Scope Authentication Selector](#) on page 214).

Requested AuthN Context Authentication Selector

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 (see [Browser-based SSO](#) on page 32 in the “Supported Standards” chapter of the *Get started with PingFederate* guide). Configured authentication sources are mapped either to SAML-specified contexts or any ad-hoc context agreed upon between the IdP and SP partners (see [Configure the HTTP Header Authentication Selector](#) on page 212).



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.

Commercial adapters and selectors

Ping Identity regularly develops and maintains integration kits, including adapters, to work with applications and leading identity management systems. Available kits may be downloaded from the Ping Identity [Downloads](#) website. Additional authentication selectors may also be added to the download site periodically; contact sales@pingidentity.com if you are looking for 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. See the PingFederate [SDK developer's guide](#) on page 559.

Hierarchical plug-in configurations

PingFederate allows you to use a configuration of an adapter (as well as certain other PingFederate plug-ins) as a *parent* instance from which you can create *child* instances. You can then modify the inherited configuration for the child instances as needed. This feature provides easier management of adapter settings in cases where only small changes to an existing adapter (or plug-in) configuration need to be made for a particular use case.

For example, different SP-connection adapter instances might have their own IdP logon 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.

 **Tip:** You can also override adapter instances as part of mapping them into either SP or IdP connections, for cases where overridden settings may apply only to one particular connection configuration.

Any changes to a parent configuration are propagated to its child (or connection-based) configurations provided the changes are not already overridden in the derived instance.

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

- Token Translators (see [Token processors and generators](#) on page 59)
- Access Token Management instances (see [OAuth access token management](#) on page 245)
- Password Credential Validators (see [Manage password credential validators](#) on page 196)
- Identity Store Provisioners (see [Configure Identity Store Provisioners](#) on page 334)

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 (the assertion) containing user-identity information that a service provider (SP) can translate, or map, to local user stores. (For more information about SAML, see [Supported standards](#) on page 29 in the *Get started with PingFederate* guide.)

For browser-based SSO, PingFederate enables two modes of identity mapping between domains:


- [Account linking](#) on page 69
- [Account mapping](#) on page 70

For WS-Trust STS, account mapping is used.

Account linking

Under the standards, *account linking* can be used for browser-based 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*, may be either a unique attribute, such as an email address, or a *pseudonym* generated by the IdP to uniquely identify individual users. Pseudonyms can be used when privacy is a concern; they cannot easily be traced back to a user's identity at the partner site.

During the user's first SSO request, the 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.

 **Tip:** An SP PingFederate uses a default, HyperSQL database to handle account linking. You can use your own data store instead, as needed. For more information, see [Define an account-linking data store](#) on page 166.

Optionally, additional attributes may be sent with the name identifier. When a pseudonym is used as the account link, however, care must be taken to send only general attributes (a user's organizational role or department, for example) 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* (see [/sp/defederate.ping](#) on page 454). A user who has defederated may later elect to reassociate with a local user account.

SP affiliations

Under the SAML 2.0 specifications, an IdP can configure PingFederate to enable a group of SPs—an *SP affiliation*—to share the same persistent name identifier (see [Define SP affiliations](#) on page 327). This capability facilitates the use case in which a number of business partners have an existing relationship and sharing a single name identifier among all parties reduces the federation integration effort.

Account mapping

Account mapping (also called “*attribute mapping*”) enables an 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. The attributes used to look up the user are always “exposed”; that is, they are known to both the IdP and SP. An email address, for example, is a commonly used identifying attribute.

Account mapping can be used to achieve one-to-one mapping (individual user accounts exist on both sides of federated connection) or many-to-few (IdP users without accounts at destination sites may be mapped to guest accounts or to a role-based general account).

For browser-based 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, whereby the user is mapped 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 may be sent with the transient identifier. Again, care should be taken to preserve privacy.

Account mapping is commonly implemented in B-to-B or B-to-E use cases where it might be appropriate for the administrator to create a user lookup on behalf of the user.

About 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 IdP can pass other user attributes in an assertion (including SAML tokens for Web Services). This supplemental information can be used by the SP for several purposes. For example, attributes may be used to map and authorize the user into a specific role, with associated site permissions. In other cases, attributes may be used to customize the end application display for a more robust user experience.


The SP also has the option of incorporating additional attributes prior to creating a session for the target application. This is commonly done where the SP also maintains an account for the user and wants to pass 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 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.

Attribute contracts

An attribute contract represents an agreement between partners about user attributes sent in a SAML assertion, a JSON web tokens (JWTs), or an OpenID Connect ID token. The contract is a list of case-sensitive attribute names. Partners must configure attribute contracts to match.


 **Tip:** When privacy is required for sensitive attributes, you can configure PingFederate to mask their values in log files (see [Attribute masking](#) on page 74).

For an IdP or an OpenID Provider (OP), the attribute contract defines which attributes PingFederate sends in an assertion, a JWT, or an ID token. While this contract is fixed for all users authenticating to the partner, the values used to fulfill the contract may differ from one user to the next. The attribute contract may be fulfilled by relying on a combination of different data sources:

- The IdP adapter or STS token processor
- An IdP attribute source, which identifies the location of individual attributes in a data store
- Static text values for some attributes, or text values combined with variables
- Expressions (see [Attribute mapping expressions](#) on page 497)

For an 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. PingFederate can be configured to pass these attributes to the SP adapter or, for Web Services, to the SP token generator (see [Manage SP adapters](#) on page 331 or [Manage token generators](#) on page 412). You can also use attributes to look up additional attributes in local data stores, which may be needed to start a user session or create a local security token for Web Services (see [Adapter contracts](#) on page 72 below or [STS token contracts](#) on page 72).

The attribute contract always contains the user identifier (SAML_SUBJECT in a SAML assertion and sub in a JWT or an ID token)—the primary information used to identify the user—unless you are using account linking for browser-based 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 session-creation adapters for two separate applications, a single attribute contract can be created for the IdP connection partner. This single contract would be constructed to supply all the attributes needed by both SP adapters.

Name formats

By agreement with an SP partner, an IdP may specify a format (email, for example) associated with the SAML_SUBJECT. The SP may require this information to facilitate handling of the format.

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

For browser-based SSO connections, 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 can also define a customized selection of additional attribute formats (see [Set up an attribute contract](#) on page 292).

 **Note:** The designation of formats is not applicable to SP administrators. The information is simply available in the incoming assertion to an SP application that might need it for particular processing requirements.

For the WS-Trust IdP configuration, attribute-name formats cannot be specified. If needed, however, an administrator can use a special variable in the attribute contract to set the subject-name format (see [Define an attribute contract for IdP STS](#) on page 406). (The same variable is also available for browser-based SSO attribute contracts, but the feature is deprecated.)

STS namespaces

By agreement with an SP partner for a WS-Trust STS connection, an IdP may specify an XML namespace to be associated with an attribute (for example, to use claims-based authorization with WIF clients—see [Windows Identity Foundation clients](#) on page 60). Namespaces can be specified only for attributes of a WS-Trust IdP configuration using SAML 1.1 or SAML 1.1 for Office 365 as the default token type (see [Define an attribute contract for IdP STS](#) on page 406).

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 IdP side of a federation, adapter attributes are supplied to PingFederate by an IdP adapter (see [SSO integration kits and adapters](#) on page 67 and [Manage IdP adapters](#) on page 275).

On the SP side, adapter contract attributes are those required by an adapter to start a session with an application. At least one *adapter type* is needed for each security domain. Then an *adapter instance* must be configured for each target application. (See [Manage SP adapters](#) on page 331.)

Adapter contracts on the SP side are fulfilled using attributes from the attribute contract, possibly enhanced through other attributes looked up from local data stores. For example, if several target applications are controlled by the same security context and can receive 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 (see [Manage target session mappings](#) on page 348).

Extended adapter contract

Adapter contracts are created when an adapter type is deployed with PingFederate. When developed, these adapters are “hard-wired” to look up or set a specific set of attributes. After deployment, your attribute requirements may change. To streamline adjustment of adapter contracts, PingFederate allows an administrator to add additional attributes to the adapter instance through the administrative console. These adjustments are called *extended adapter contracts*.

STS token contracts

Similar to an adapter contract for browser-based SSO, an 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 IdP side of a federation, token-processor attributes are supplied to PingFederate (see [Token processors and generators](#) on page 59 and [Manage token processors](#) on page 401).

On the SP side, token-generator contract attributes are those required by a token generator to pass identity information from the token to the Web Service client application. At least one token generator type is needed for each security domain. Then a token generator instance must be configured for each target application (see [Manage token generators](#) on page 412). 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 (see [Manage SP token generator mappings](#) on page 415).

Extended token generator contract

Token-generator contracts are created when a token-generator type is deployed with PingFederate. When developed, these token generators are “hard-wired” to look up or set a specific set of attributes. After deployment, your attribute requirements may 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*.

Data stores

PingFederate can be configured to use local data stores to supply attributes for either the IdP's attribute contract, the SP's adapter contract, or STS token contracts (see sections above). Standard data stores may include any JDBC-accessible database or an LDAP v3-compliant directory server (see [Manage data stores](#) on page 158).

Alternatively, you can use the PingFederate Custom Source SDK to create your own driver for non-JDBC/LDAP data stores—including, for example, flat files or SOAP-connected databases (see the PingFederate [SDK developer's guide](#) on page 559).

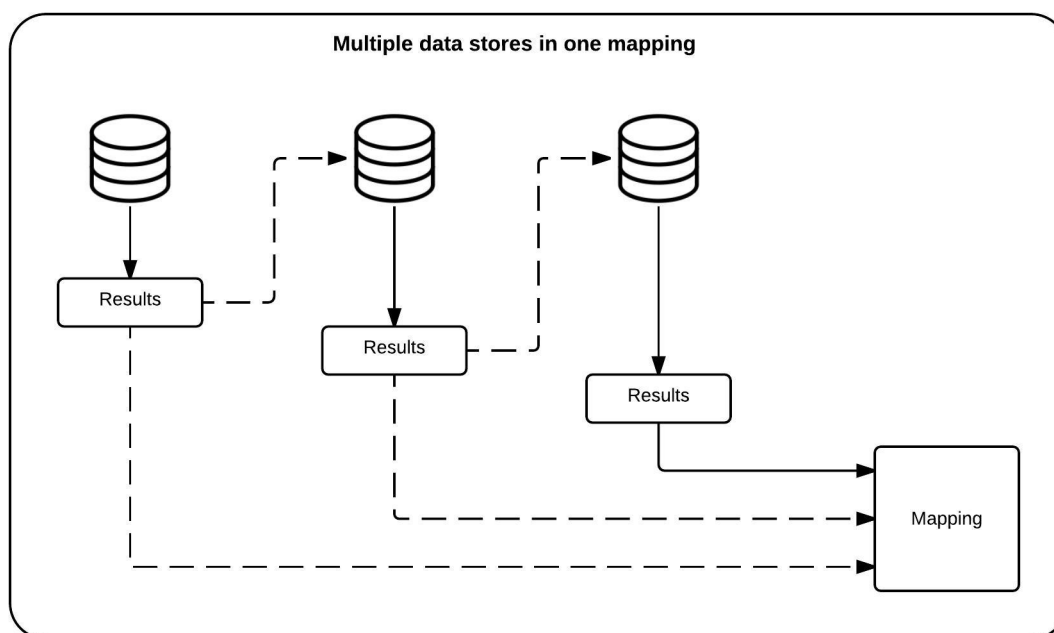
Data stores can be used across multiple connections.

Multiple data source attribute mapping

When querying data stores for attributes to help fulfill a contract on the IdP side, PingFederate® can be configured to use more than one attribute source.

Multiple data stores in one mapping

The PingFederate IdP server supports separate data stores to look up attributes for a single mapping. For example, you can query multiple data stores for values and map those values in one mapping, or query a data store for a value and use that value as input for subsequent queries of other data stores.



If a data store 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 data store 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) cannot be fulfilled, the request fails.

Multiple data stores in one mapping is available for Browser SSO and WS-Trust STS on the IdP side, authentication policy contract to SP adapter mappings, and the following OAuth workflows:

- OpenID Connect policies (the user-attributes mapping process)
- Resource-owner credential mappings
- IdP adapter mappings
- Access token mappings

Multiple mappings and failsafe mapping

For Browser SSO and WS-Trust STS on the IdP side, PingFederate also supports separate search parameters for each data store and for "fall-through" searches. For example, you can add the same data store more than once, using different search queries for each instance, or you can search different data stores successively. If any search fails to find a user in the specified attribute source, the next search is executed until a match is found. A failsafe attribute source can also be configured, providing a default set of attributes from the IdP adapter and using text values.



Attribute masking

At runtime PingFederate® logs user attributes (see [PingFederate log files](#) on page 95). To preserve user privacy, you may wish to mask the values of logged attributes.

PingFederate provides this masking capability at all points where the server logs attributes. These points include:

- Data-store lookup at either the IdP or SP site (see [Manage data stores](#) on page 158).
- Retrieval of attributes from an IdP adapter or token processor (see [Set pseudonym and masking options](#) on page 276 and [Set attribute masking](#) on page 404).
- SP-server processing of incoming attributes based on the SSO attribute contract, (see [Define an attribute contract](#) on page 348).

Note that 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 (see [Account linking](#) on page 69).

- SP-server processing of incoming attributes in response to an Attribute Request under XASP (see [Configure security policy for Attribute Query](#) on page 366).

For information about XASP, see [Attribute Query and XASP](#) on page 43 in the “Supported Standards” chapter of the *Get started with PingFederate* guide.

- ! **Important:** Many adapter implementations, as well as other product extensions, may independently write unmasked attribute values to the PingFederate server log. These implementations are beyond the control of PingFederate. If sensitive attribute values are a concern when using such a component, a system administrator can adjust the component's logging threshold in `log4j2.xml` to prevent the recording of attributes (see [PingFederate log files](#) on page 95).

About token authorization

PingFederate provides an optional configuration to evaluate user attributes, as well as other runtime variables (such as authentication context), for authorization purposes. This feature, known as *token authorization*, provides a way for administrators to extend access policy directly to many areas, such as Browser SSO, STS, and OAuth events, by conditionally allowing or disallowing the issuance of relevant security tokens (for example, SAML assertions, STS tokens, OAuth access tokens, or session cookies). The option is also available for extending authorization policy to attribute-query responses, IdP adapter contracts, and authentication policy contracts.

Administrators can configure token authorization using Issuance Criteria screens immediately following the configuration of attribute mapping at all applicable points in the administrative console (see [Specify issuance criteria](#) on page 299, as an example).

Issuance criteria


The token-authorization configuration consists of one or more rules that evaluate attribute values against selected conditions (see [Single and multi-value conditions](#) on page 75 below). You can choose from among several sources for the attributes, depending on the type of configuration that contains the token-authorization setup. The sources always consist of all of those available for attribute mapping, including data stores (when configured) and runtime information related to the context of an event. In addition, values of mapped attributes are available to provide access to any plain-text mappings or the runtime results of any expression mappings (see [Attribute mapping expressions](#) on page 497).

Tip: 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 OGNL expressions using an Advanced Criteria feature. For more information, see [Attribute mapping expressions](#) on page 497 and [Sample OGNL expressions](#) on page 498.

Note: When authorization fails and a transaction is halted, a configurable Error Result code is passed back up through the system, 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. For more information, see Error Result descriptions in sections of this document covering specific token-authorization configurations.

Single and multi-value conditions

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



Use one of the first six choices only for attributes consisting of a single value.

Note: Use one of the *multi-value* conditions when you want PingFederate to verify that an attribute contains or does not contain the specified value in its attribute list. Using a single-value condition when an attribute has multiple values causes the condition to evaluate consistently to false.

Security infrastructure

This section describes the PingFederate security infrastructure that supports encrypted messaging, certificates, and digital signing. These functions are integrated into PingFederate's configuration screens 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. Digital signatures are used in both SAML (including STS tokens) and WS-Federation electronic documents.

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

Message signing

Certificates contain information about the owner of the certificate along with a public key. Applying a digital signature creates and encrypts a hash from the message you are signing, using your private key. PingFederate provides a choice of signature encryption algorithms when a stronger algorithm is required.

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

Note: Digital signatures do not encrypt the contents of a message; XML encryption is used for this purpose.

The certificate should be signed by a Certificate Authority (recommended), but it can be self-signed or signed by an untrusted third party. After generating a keypair and a self-signed certificate, you can use PingFederate to create a Certificate Signing Request (CSR) and send it to a CA for signing. After the CA has generated a Certificate Signing Response, you can import it into PingFederate's certificate management system. (The CA's certificate must be in PingFederate's trusted store or in the Java runtime `cacerts` store.)

PingFederate enables signing and validation of requests and responses. In addition, PingFederate provides for certificate generation, import and export functionality, CSR generation, and application of digital signatures. You can create reusable global signing certificates across your federated connection base and import signature verification certificates for each partner (see [Manage digital signing certificates and decryption keys](#) on page 184).



Note: 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, both when they are initially imported and at runtime when they are used to encrypt, decrypt, and digitally sign or verify assertions.

PingFederate can also check to see whether a certificate has been revoked, using either Certificate Revocation Lists (CRLs) or the Online Certificate Status Protocol (OCSP). Depending on the content of the certificate in question and your requirements, the server will perform either of these checks during SSO or 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 is acting as the client making an HTTPS request to a separate server

If a certificate is expired or revoked, the associated SSO or SLO transaction fails at runtime and an error is written to the transaction log. In the administrative console, an expired or revoked certificate is identified as such in the Status column of its respective Certificate Management list.

CRL Revocation Checking

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

No setup is needed in the administrative console to enable CRL checking. PingFederate automatically checks CRLs if all of the following conditions are met:

- The certificate contains the URL where the CA maintains its CRL.
- The URL is accessible.
- The returned CRL is signed and the signature verified.
- CRL validation is not explicitly disabled as a failover option in the OCSP setup.

OCSP Revocation Checking

OCSP was developed as an alternative to CRL validation and provides a more centralized and potentially more reliable means of checking certificate status. In this scenario, an OCSP Responder URL is normally embedded in the incoming certificate (a configured default URL may be used, alternatively). The URL, maintained by the issuing CA, is used 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 revocation checking is handled by the Responder server, which returns the certificate status.

A PingFederate administrator can enable and configure OCSP processing in the administrative console. The protocol may be used 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 [Configure certificate revocation](#) on page 191.

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, the protocol specifications require signatures—for example, all SAML STS tokens and all SSO assertions sent across the POST binding must be signed. (These requirements are enforced by the PingFederate administrative console and the runtime protocol engine.) Other uses of the digital signatures are optional between partners and enforced if specified for a partner connection.

If a digital-signing certificate is not issued by a trusted CA (that is, “self-signed”), then the signing partner must send the public-key certificate out-of-band (for example, via email) to the partner. The partner must import the certificate into PingFederate when configuring a connection to the signing partner for SSO/SLO or STS.

If the certificate is signed by a trusted CA and the signing partner chooses to embed the certificate in all signed messages, then the verifying partner can elect to use the embedded certificate for signature verification, after validating it against the Subject DN of the original certificate. The public-key certificate may or may not be sent out-of-band (just the Subject DN is required).

 **Tip:** PingFederate can extract the Subject DN from the certificate, when available, during the signature-verification configuration.

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. Note that for each trust model, PingFederate always verifies that the certificate is current and that the signature in the message can be verified using 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 the anchored trust model is chosen, 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

When this option is chosen, incoming signatures are verified exclusively using a specific certificate imported for a connection into PingFederate (or a secondary, backup certificate when specified). The certificate may be self-signed or issued by a trusted CA. The certificate chain, if any, is not verified. However, revocation checking, when enabled, is performed up any existing chain as far as available.

Secure sockets layer

SSL certificates signed by a CA can be used to identify one or both ends of the federation. SSL/TLS provides an encrypted connection between the two parties in which the content of a message is not exposed, thus ensuring 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 IdP site. On the SP side, the Artifact Resolution Service should also use SSL/TLS. Optionally, SSL/TLS may also be used to secure communication between internal data stores and PingFederate and between the PingFederate STS and Web Service client or provider applications.

Authentication

Three methods of authentication, described below, are available for use with PingFederate for browser-based SSO to authenticate connection partners making SOAP requests. For SOAP authentication by STS clients, a separate option using either or both of the first two methods, may be configured (the third method, digital signing, is automatically required). The selection of one or more method(s) must be agreed upon between partners and synchronized 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's certificate out-of-band (see [Manage SSL client keys and certificates](#) on page 182).



Important: The SSL/TLS server-client handshake involves negotiating cipher suites to be used for encryption/decryption on each side of a secured transaction. PingFederate supports only stronger cipher suites; to enhance security, weaker cipher suites are commented out of two configuration files located in `<pf_install>/server/default/data/config-store:`

- `com.pingidentity.crypto.SunJCEManager.xml`
- `com.pingidentity.crypto.NcipherJCEManager.xml`

To ensure the most secure transactions, we recommend that administrators retain this cipher-suite configuration.

Due to import control restrictions, the standard Java Runtime Environment (JRE) distribution supports strong but not unlimited encryption. For this reason, the strongest cipher suites in the same configuration files are also commented out. To use the strongest encryption, when permissible, remove the comments from the AES 256 cipher suites and download and install the appropriate version of “Java Cryptography Extension (JCE) Unlimited Strength Jurisdiction Policy Files” from the [Oracle download website](http://www.oracle.com/technetwork/java/javase/downloads/jce8-download-2133166.html) (www.oracle.com/technetwork/java/javase/downloads/jce8-download-2133166.html).

- Digital Signatures: partners sign the XML message transmitted via the SSL/TLS connection. Signatures are verified by the receiver based upon the certificate(s) configured for that connection. Each partner should import the other's certificate(s) out-of-band (see [Manage digital signing certificates and decryption keys](#) on page 184).

Verifying trusted certificates

PingFederate validates the trust of all certificates. A certificate is trusted if the certificate of its issuer is in PingFederate's trusted certificate store. The root certificate of the CA, by which a certificate is issued, must be imported into PingFederate's trusted certificate store or contained in the Java runtime `cacerts` store.

Encryption

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

For SAML 2.0 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.

Using Auto-Connect

PingFederate allows organizations to provide secure SSO on the fly—that is, without the need for configuring partner-specific, browser-based SSO connection parameters. This feature—*Auto-Connect*[™]—extends SAML 2.0 SP-initiated SSO or SLO and metadata specifications to enable deployments to retrieve partner connection information securely on an as-needed basis. (For information about SAML 2.0, see [Supported Standards](#) in the *Get started with PingFederate* guide.)

The feature is especially useful to an SP who wants to provide SSO capability to more than one partner. A SaaS provider, for example, can provide SSO to innumerable clients without specifying redundant connection information for each one. Auto-Connect can also help an enterprise, acting as an IdP, to provide easily scalable SSO for multiple outsourced services.

For either an IdP or SP PingFederate server, you can implement Auto-Connect for any number of partners by configuring a common initial setup and a list of domain names. For an IdP, the domain-name list contains SP partners from whom your site will accept Auto-Connect authentication requests. For an SP, the list contains IdP-partner domains to which your site can send authentication requests and receive SSO assertions.

For information about configuring Auto-Connect for your federation partners, see [Configure SP Auto-Connect](#) on page 329 or [Configure IdP Auto-Connect](#) on page 398.

Providing metadata

You enable Auto-Connect as part of Server Settings from the Main Menu (see [Choose roles and protocols](#) on page 151). Once Auto-Connect is enabled and the initial setup is fully configured and activated, partners can retrieve your connection metadata via HTTP. At runtime, Auto-Connect deployments at partner sites use the endpoints provided in the metadata to interact with your server and complete SSO or SLO processing.

The metadata, which follows SAML 2.0 specifications, must be signed, and the validity of the data is time-limited (see [Auto-Connect security model](#) on page 80 and [Configure metadata lifetime](#) on page 157).

Runtime processing

Auto-Connect runtime processing starts when a user tries to reach a protected SP resource. The process depends on SP Web-application functionality that determines the user's IdP domain (for example, from a submitted email address) and passes it to the SP PingFederate server in the SSO request.

This illustration and the accompanying processing steps describe the complete SSO processing flow:

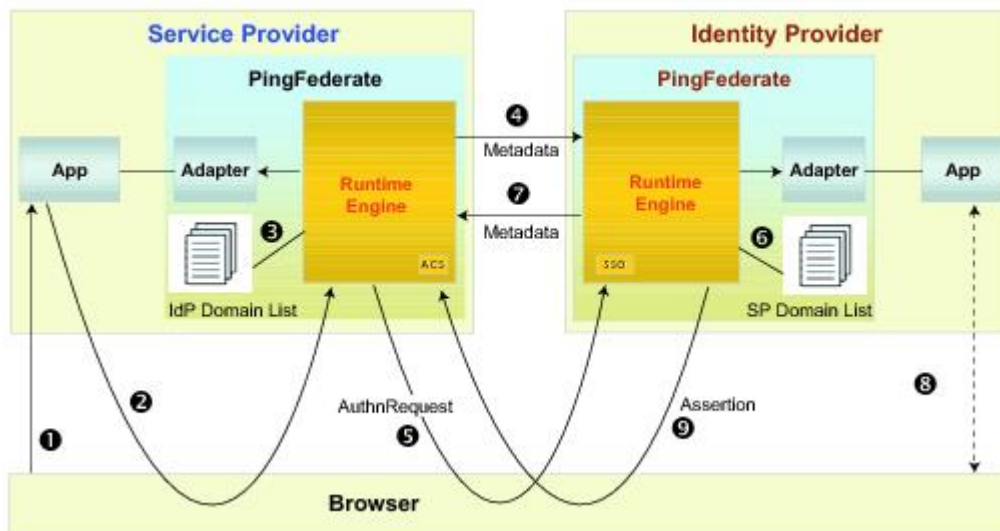


Figure 17: Auto-Connect Processing Flow

Processing steps

1. User sends a logon request with an email address to an SP application. For example:

john@example.com

- The application parses the email address and sends a request to PingFederate. For example:

```
https://www.example.org:9031/sp/startSSO.ping/?Domain=example.com
```

- The SP PingFederate server looks up the domain in a list of domain names allowed to use Auto-Connect.
- If the domain is in the list, the SP retrieves connection metadata from the IdP's public endpoint.

By default, PingFederate looks for the metadata by prepending `http://saml` to the domain. For example:

```
http://saml.example.com
```

This default location can be changed, if necessary, in the Allowed Domains lists configured in the PingFederate administrative console.

- After validating the metadata (see [Auto-Connect security model](#) on page 80), the SP sends an authentication request to the IdP's SSO service.
- If the request `<Issuer>` is not among the IdP's static-connection partners, the IdP PingFederate server looks for the issuer's domain name in the list of domains allowed to use Auto-Connect.
- The IdP retrieves the SP's metadata via its public endpoint and verifies the metadata signature.

The process is the same as that used by the SP in [Step 4](#).

- The IdP requests user authentication via the configured adapter instance.
- Once the user is authenticated, the IdP returns a signed SAML assertion to the SP's Assertion Consumer Service (ACS) endpoint.
- (Not shown) The SP logs the user on to the requested resource via the configured SP adapter.

Auto-Connect security model

Auto-Connect processing requires digital signatures to ensure the authenticity of the published metadata as well as all subsequent SSO or SLO requests and responses. The certificate used to sign the metadata is included in the metadata, and all certificates must be signed by a trusted Certificate Authority; thus, partners need not exchange certificates out of band.

In addition to validating certificates, the PingFederate runtime server compares the partner certificate with the entity ID (the "Issuer") found in the SAML message. Then the server matches the entity ID against the configured list of allowed Auto-Connect domains.

This diagram illustrates the security validation process:

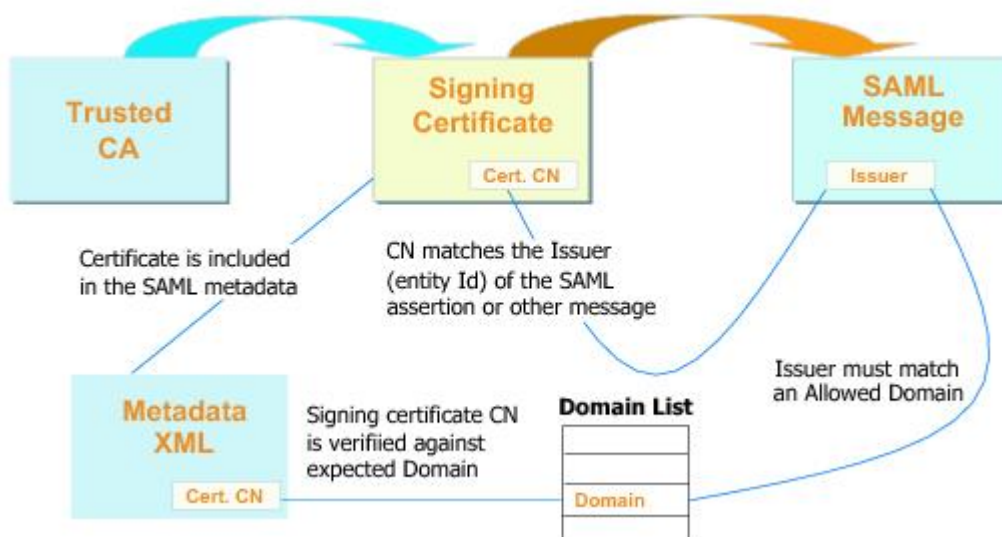


Figure 18: Auto-Connect Security Model

Note that the diagram assumes that the same certificate is used for signing both the metadata and the runtime SAML messages. This is convenient, but not required.

User provisioning

PingFederate provides cross-domain user provisioning and account management. User provisioning is an important aspect of identity federation. Often 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 may be configured either in conjunction with partner SSO connections or separately:

- At an IdP site, you can automatically provision and maintain user accounts at service-provider sites that have implemented the System for Cross-domain Identity Management (SCIM), or at selected SaaS providers (see the next section, *Outbound provisioning for IdPs* on page 81).

For information about SCIM, please refer to www.simplecloud.info.

- When PingFederate is configured as an SP, you can provision and manage 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 (see *Provisioning for SPs* on page 82).

Outbound provisioning for IdPs

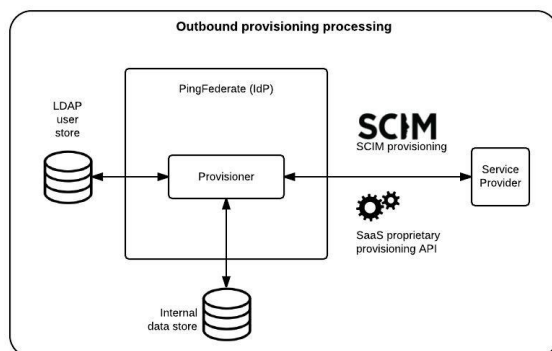
User provisioning is an important aspect of identity federation. Often when organizations enable SSO for their users, they must ensure that some form of account synchronization is in place. Automated user provisioning features within PingFederate® free administrators from having to devise a manual strategy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

For IdP sites, PingFederate provides built-in automated provisioning and user-account management to SCIM-enabled services providers and to selected SaaS providers, via their proprietary provisioning APIs.

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

- **Tip:** Support for provisioning for SaaS applications, including quick-connection templates to expedite the configuration effort, is available separately. Contact sales@pingidentity.com for more information.

When outbound provisioning is 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, and Oracle Directory Server; templates are used to preconfigure many provisioning settings. Although these are the only data stores formally tested and supported, other LDAP data stores will likely work as well.

Internal data store

Tested internal data stores used for synchronization include HyperSQL, Microsoft SQL Server, Oracle Databases, Oracle MySQL, and PostgreSQL. A demonstration-only, embedded HyperSQL database is installed by default. Again, any relational database may be used; scripts are provided to aid setup.

Provisioning for SPs

User provisioning is an important aspect of identity federation. Often 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:

- SCIM inbound provisioning for SCIM requests coming either from inside or outside an organization
- Just-in-time provisioning for creating and updating user accounts based on information contained in SSO tokens

Inbound provisioning

SCIM inbound provisioning provides support for incoming SCIM messages containing requests to create, read, update, or delete (or deactivate) user and group records in Microsoft Active Directory data stores or custom user stores via the Identity Store Provisioners. PingFederate supports SCIM attributes in the core schema and custom attributes through a schema extension. An administrator can configure this provisioning feature by itself or in conjunction with an SSO or other connection types.

In effect, inbound provisioning provides an organization with a dedicated SCIM service provider, which can route user-management requests to an organization's centralized user store. The requests may originate from trusted applications within an organization (for example, a human-resources on-boarding SaaS product) or from trusted partner IdPs.

For setup information, see [Configure SCIM inbound provisioning](#) on page 372. To integrate Inbound Provisioning with custom user stores, see [Configure Identity Store Provisioners](#) on page 334. For application-development information about using PingFederate endpoints for SCIM provisioning, see [SCIM inbound provisioning endpoints](#) on page 455.

Just-in-time provisioning

At an SP site, PingFederate can create and update local user accounts in an external LDAP directory or Microsoft SQL Server as part of SSO processing—*Just-in-time (JIT) provisioning* (also formerly known as Express Provisioning). This feature allows SPs to maintain accounts for users who authenticate via IdP partners without having to provision accounts manually, when local accounts are required.

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, assertion attributes can be supplemented with user attributes returned from an Attribute Query.

PingFederate can also update existing user accounts based on assertions. When this option is enabled, PingFederate can add or overwrite attributes for a local user account each time SSO for a user is processed.



Note: Note that once user attributes are provisioned, they cannot be removed using JIT provisioning. Where deprovisioning is required, we recommend using SCIM inbound provisioning.

For information about enabling JIT Provisioning, see [Choose IdP connection options](#) on page 342. For configuration information, see [Configure just-in-time provisioning](#) on page 366.

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.



Tip: Extensive coordination and configuration may be avoided by using Auto-Connect (see [Using Auto-Connect](#) on page 79).

Standards and Specifications

Choose which federation protocol(s) your deployment will support. For SAML SSO configurations, decide which profiles and bindings will be used. (See [Supported Standards](#) chapter of the *Get started with PingFederate* guide.)

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). (See [Security infrastructure](#) on page 75.)

Also, if a stronger signature algorithm is required, determine what RSA algorithm will be used for signing. (The optional algorithm selection is available throughout the administrative console, where signing certificates are specified for various uses.)

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 data stores 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 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 (see [Trust models](#) on page 77).


Deployment

Decide how PingFederate fits into your existing network. Also, determine whether high-availability, failover options, or both, are required (see the PingFederate [Server clustering guide](#) on page 533).

Federation Server Identification

Determine how you and your partner(s) will identify your respective federation deployments. Under federation standards, both the sender (IdP) and the receiver (SP) of an assertion must be uniquely identified within the identity federation (see [Configuration data exchange](#) on page 85).

With PingFederate, you define a unique ID for each supported protocol (see [Specify federation information](#) on page 152). Optionally, you can also use a list of multiple *Virtual Server IDs* on a connection-by-connection basis (see [Multiple virtual server IDs](#) on page 84).

 **Tip:** PingFederate also provides for *virtual host names*, which differ from virtual server IDs (but are not mutually exclusive); they are intended to be used when your network configuration is such that you receive federation messages under more than one domain name (see [Configure virtual host names](#) on page 115).

Server Clock Synchronization

Ensure that both the SP and IdP server clocks are synchronized. SAML messages and 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 data store that contains user data when needed: LDAP, JDBC, or Custom (see [Data stores](#) on page 72).

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.

(See [SSO integration kits and adapters](#) on page 67 and [Token processors and generators](#) on page 59.)

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 95.)

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 (see [Identity mapping](#) on page 69).

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 71.)

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. (If your partner is also using PingFederate or supports standards permitting runtime metadata exchange, the process can be even simpler—see [Using Auto-Connect](#) on page 79.)

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 SP connection to access multiple subdomain accounts in Microsoft Office 365.

Suppose both the marketing and the engineering departments of contoso.com (the 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.

In this scenario, the PingFederate IdP server can be configured to include both marketing.contoso.com and engineering.contoso.com as the virtual server IDs in the Office 365 SP connection. Each virtual server ID has its own set of protocol endpoints, which can be obtained in the connection metadata (see [Export metadata to an XML file](#) on page 116 and [System-services endpoints](#) on page 460 for more information).

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 issuer being 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 (see [Identify the SP](#) on page 288) or as an SP configuring an IdP connection (see [Identify the partner](#) on page 344) for both Browser SSO Profiles and WS-Trust STS (for access to identity-enabled Web Services).

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 default virtual server ID is used when no virtual server ID information is included in a request (see [IdP endpoints](#) on page 448 as an IdP or [SP endpoints](#) on page 451 as an SP).

In a connection with multiple virtual server IDs, you can optionally restrict each adapter added to the connection to certain virtual server IDs to enhance the end-user experience (see [Restrict an authentication source to certain virtual server IDs](#) on page 296, [Restrict a target session to certain virtual server IDs](#) on page 350 and [Restrict a target session to certain virtual server IDs](#) on page 350).

 **Tip:** You can also restrict each token processor or token generator added to a WS-Trust STS SP connection or IdP connection, (see [Restrict a token processor to certain virtual server IDs](#) on page 408 or [Restrict a token generator to certain virtual server IDs](#) on page 416).

 **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 [Specify issuance criteria](#) on page 299 as an IdP or [Identify issuance criteria](#) on page 352 as an SP.

Configuration data exchange

If your partner's deployment does not produce or consume a metadata file that conforms to SAML metadata specifications, you may need to exchange connection information manually. The following sections list some common configuration details that must be exchanged if metadata files are not used. (These lists are not exhaustive.)

IdP to 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 are configuring):

- **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 (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 Service URL**—The destination of SLO request messages.
- **Single Sign-On 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 are configuring):

- **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 (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

Many settings must be mutually set by the parties. This information might include such items as:

- **Attributes**—User information that will be sent in an assertion, if any (see [About attributes](#) on page 70).
- **Signing certificates**—The SAML and WS-Federation protocols specify a number of conditions under which digital signatures are either required or optional (these conditions are built into the PingFederate connection-setup screens).
- **SOAP connection type and authentication style**—For SAML connections using the back channel (using the artifact binding, for example), 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.

Federation hub use cases

PingFederate® can be configured 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 can bridge browser-based SSO between identity providers and service providers. It stands in the middle of the SSO and SLO flow, acting as the SP for the identity providers and as the IdP for the service providers. The four use cases are:


- [Bridging an IdP to an SP](#) on page 86
- [Bridging an IdP to multiple SPs](#) on page 87
- [Bridging multiple IdPs to an SP](#) on page 87
- [Bridging multiple IdPs to multiple SPs](#) on page 88

PingFederate also supports protocol translation among SAML 1.0, 1.1, 2.0, OpenID Connect (for identity providers), 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 can be deployed 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

In this use case, PingFederate is bridging SSO and SLO transactions between an identity provider and a service provider. For example, you may have a legacy IdP system that is only capable of sending SAML 1.1 assertions via POST. Your service provider however requires SAML 2.0 assertions via the artifact binding. With federation hub, you can configure PingFederate to consume inbound SAML 1.1 assertions (by POST), translate them to SAML 2.0 assertions, and send them via the artifact binding to the service provider.




To address this use case:

1. Enable both the IdP and the SP roles with the applicable protocols on the **Server Configuration > Server Settings > Roles & Protocols** screen.

2. Create a contract to bridge the attributes between the identity provider and the service provider (see [Federation hub and authentication policy contracts](#) on page 89).
3. Create an IdP connection between the identity provider and PingFederate (the federation hub as the SP) and add to the IdP connection the applicable authentication policy contract(s) on the **Target Session Mapping** screen.
4. Create an SP connection between PingFederate® (the federation hub as the IdP) and the service provider and add to the SP connection the corresponding authentication policy contract on the **Authentication Source Mapping** screen.
5. Work with the identity provider to connect to PingFederate (the federation hub) as the SP.
6. Work with the service provider to connect to PingFederate (the federation hub) as the IdP.

Bridging an IdP to multiple SPs

In this use case, PingFederate is bridging SSO and SLO transactions between an identity provider and multiple service providers. For example, your company wants to route federation requests from a recently acquired subsidiary through its federation infrastructure. With PingFederate, you can multiplex one IdP connection to multiple SP connections to the desired service providers. The federation hub consumes assertions from the subsidiary and creates new assertions to the respective service providers.




To address this use case:

1. Enable both the IdP and the SP roles with the applicable protocols on the **Server Configuration > Server Settings > Roles & Protocols** screen.
2. For each service provider, create a contract to the identity provider (see [Federation hub and authentication policy contracts](#) on page 89). Multiple contracts are likely required, because each service provider may require a unique set of attributes.
3. Create an IdP connection between the identity provider and PingFederate (the federation hub as the SP) and add to the IdP connection the applicable authentication policy contract(s) on the **Target Session Mapping** screen.
4. For each service provider, create an SP connection between PingFederate® (the federation hub as the IdP) and the service provider and add to the SP connection the corresponding authentication policy contract on the **Authentication Source Mapping** screen.
5. For each service provider supporting the SAML IdP-initiated SSO profile, map the expected target resources to the corresponding SP connections on the **SP Configuration > Target URL Mapping** screen.
6. Work with the identity provider to connect to PingFederate (the federation hub as the SP).
7. Work with each service provider to connect to PingFederate (the federation hub as the IdP).

Bridging multiple IdPs to an SP

In this use case, PingFederate is bridging SSO and SLO transactions between multiple identity providers and a service provider. 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.



To address this use case:

1. Enable both the IdP and the SP roles with the applicable protocols on the **Server Configuration > Server Settings > Roles & Protocols** screen.
2. Create a contract to bridge the attributes between the identity providers and the service provider (see [Federation hub and authentication policy contracts](#) on page 89). You likely need only one contract unless the service provider requires a different set of attributes from each identity provider.
3. For each identity provider, create an IdP connection between the identity provider and PingFederate (the federation hub as the SP) and add to the IdP connection the applicable authentication policy contract(s) on the **Target Session Mapping** screen.
4. On the **Selectors** screen, configure an authentication selector to map each identity provider to the corresponding IdP connection.
5. Create an SP connection between PingFederate® (the federation hub as the IdP) and the service provider and add to the SP connection the corresponding authentication policy contract on the **Authentication Source Mapping** screen.



Important: PingFederate includes the Entity ID of the original identity provider (Authenticating Authority) in SAML 2.0 assertions so that the service provider can determine the original issuer of the assertions. This is especially important when bridging multiple identity providers to one service provider—the service provider 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** screen and then map **Context: Authenticating Authority** as the attribute value on the **Attribute Contract Fulfillment** screen.

For information about Authenticating Authority, see section [2.7.2.2 Element <AuthnContext>](#) in the SAML 2.0 specification (<https://docs.oasis-open.org/security/saml/v2.0/saml-core-2.0-os.pdf>)




Note: If the service provider does not take action based on Authenticating Authority, depending on the attributes from the identity providers, you may define validation rules on the **Issuance Criteria** screen to protect against user impersonation between IdPs.

6. Work with each identity provider to connect to PingFederate (the federation hub as the SP).
7. Work with the service provider to connect to PingFederate (the federation hub as the IdP).

Bridging multiple IdPs to multiple SPs

This use case is a combination of [Bridging an IdP to multiple SPs](#) on page 87 and [Bridging multiple IdPs to an SP](#) on page 87.



To address this use case:

1. Enable both the IdP and the SP roles with the applicable protocols on the **Server Configuration > Server Settings > Roles & Protocols** screen.
2. Create multiple contracts to bridge the attributes between the identity providers and the service providers (see [Federation hub and authentication policy contracts](#) on page 89).
3. For each identity provider, create an IdP connection between the identity provider and PingFederate (the federation hub as the SP) and add to the IdP connection the applicable authentication policy contract(s) on the **Target Session Mapping** screen.
4. On the **Selectors** screen, configure an authentication selector to map each identity provider to the corresponding IdP connection.
5. For each service provider, create an SP connection between PingFederate® (the federation hub as the IdP) and the service provider and add to the SP connection the corresponding authentication policy contract on the **Authentication Source Mapping** screen.



Important: PingFederate includes the Entity ID of the original identity provider (Authenticating Authority) in SAML 2.0 assertions so that the service provider can determine the original issuer of the assertions. This is especially important when bridging multiple identity providers to one service provider—the service provider 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** screen and then map **Context: Authenticating Authority** as the attribute value on the **Attribute Contract Fulfillment** screen.

For information about Authenticating Authority, see section [2.7.2.2 Element <AuthnContext>](#) in the SAML 2.0 specification (<https://docs.oasis-open.org/security/saml/v2.0/saml-core-2.0-os.pdf>)



Note: If the service provider does not take action based on Authenticating Authority, depending on the attributes from the identity providers, you may define validation rules on the **Issuance Criteria** screen to protect against user impersonation between IdPs.

6. For each service provider supporting the SAML IdP-initiated SSO profile, map the expected target resources to the corresponding SP connections on the **SP Configuration > Target URL Mapping** screen.
7. Work with each identity provider to connect to the federation hub (as the SP).
8. Work with each service provider to connect to the federation hub (as the IdP).

Federation hub and authentication policy contracts

PingFederate uses two connections to bridge an identity provider to a service provider:

- 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

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

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 have the option to configure the IdP connections, SP connections, or both, to look up values from multiple data store instances.

When bridging one identity provider to one service provider, you need to create one authentication policy contract and associate the contract with both the IdP connection and the SP connection.

When bridging one identity provider to multiple service providers, you need to create an authentication policy contract per service provider because each service provider likely requires a different set of attributes. Map all the authentication policy contracts into the IdP connection. Add the respective authentication policy contract to each SP connection to the service provider.

When bridging multiple identity providers to one service provider, you likely need only one contract unless the service provider requires a different set of attributes from each identity provider. Add the authentication policy contract to the SP connection and the applicable IdP connections.

Authentication policy contracts are managed using the **Policy Contracts** wizard. You can access it from the **IdP Configuration** or **SP Configuration** menu.

Federation hub and virtual server IDs

PingFederate uses two connections to bridge an identity provider to a service provider:

- 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 identity provider through the IdP connection and generates new assertions to the service provider via the SP connection.

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

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), for SP-initiated SSO, if the service provider sends AuthnRequest messages to the virtual server ID specific endpoint, PingFederate retains this information automatically. When the identity provider returns the corresponding assertions to PingFederate (the federation hub as the SP), PingFederate retrieves the preserved information and uses that specific virtual server ID as the issuer in the assertions it sends to the service provider. For IdP-initiated SSO, the issuer of the assertions (to the service provider) is the default Virtual Server ID.

System administration

This chapter describes general administrative functions for PingFederate, including:

- [PingFederate log files](#) on page 95
- [Export metadata to an XML file](#) on page 116
- [Sign XML files](#) on page 118
- [Configuration archive](#) on page 118
- [Account management](#) on page 108
- [Alternative console authentication](#) on page 112
- [Manage email configuration](#) on page 114
- [Configure virtual host names](#) on page 115
- [Configure PingFederate properties](#) on page 93
- [Manage PingFederate license](#) on page 91
- [Automating configuration migration](#) on page 120

- [Outbound provisioning CLI](#) on page 124



Note: The information in this chapter is presented from the viewpoint of an administrative user with “Admin” permissions (see [Account management](#) on page 108).

Manage PingFederate license

When you access the PingFederate® administrative console for the first time, you are prompted to import a license file. This applies to new and upgraded PingFederate instances.

Depending on your licensing agreement, your PingFederate license may 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.

To review the existing license key:

1. Start PingFederate and sign on to the administrative console.
2. Point to the user icon in the upper-right corner of the administrative console.
3. Click **About** in the drop-down menu.

The license summary is displayed in a pop-up browser window.

Note that if an expiration date is specified, the license expires at the beginning of the day.

Request a new license key

1. Go to the Ping Identity [licensing](#) website.
2. Sign on, provide the required information, and then submit your request.

You shall hear from our licensing team as the team processes your request.

Install a license key on a new or upgraded PingFederate server

1. Start PingFederate and sign on to the administrative console.
2. On the **Import License** screen, click **Choose file** to select the license file, and then click **Import**.

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.

Once 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 the **Server Configuration > Cluster Management** screen and click **Replicate Configuration**.



Note: Starting with PingFederate 8.2, you must use the **Replicate Configuration** workflow 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.



Note: 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 referred to as *the secondary license file* in the previous versions of PingFederate. The `.pingfederate.lic` file, if any, is ignored.

Install a replacement license key

You may replace your existing PingFederate license key by importing a new license file in the **Server Configuration > License Management** screen.

1. Start PingFederate and sign on to the administrative console.
2. Go to the **Server Configuration > License Management** screen.
3. On the **License Management** screen, click **Choose File** to select the license file, and then click **Import**.


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 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, you are advised and prompted on whether to continue.

Once 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.

4. If you have a clustered PingFederate environment, go to the **Server Configuration > Cluster Management** screen and click **Replicate Configuration**.

 **Note:** Starting with PingFederate 8.2, you must use the **Replicate Configuration** workflow 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 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.
```

 **Tip:** 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, and then click **Replicate Configuration** on the **Cluster Management** screen again.

If the issue date of the new license is more recent than 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:** 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 referred to as *the secondary license file* in the previous versions of PingFederate. The `.pingfederate.lic` file, if any, is ignored.

Configure notification for licensing events

If your PingFederate® license has an expiration date, you may configure PingFederate to notify the administrators ahead of the license expiration date to minimize potential service disruptions.

1. Go to the **Server Configuration > Server Settings > Runtime Notifications** screen.
2. Select the **Notification for Servers Licensing Events** check box.

Note that this check box appears only if your PingFederate license has an expiration date.



3. Enter an email address, and then click **Save**.


If you have not yet configured email server settings, click **Email Server Settings** to provide the required information. When you complete the email server configuration, the administrative console brings you back to the **Runtime Notifications** screen. Click **Save**.



Configure PingFederate properties

The default administrative-console and runtime behavior of PingFederate® is controlled in part by configuration properties contained in the file `run.properties`. The most common properties are documented in the following table. For the rest of the properties including various cookie-encoding options, refer to the file itself.

Table 2: PingFederate Configuration Properties

Property	Description
<code>pf.admin.https.port</code>	Defines the port on which the PingFederate administrative console runs. The default value is 9999.
<code>pf.console.bind.address</code>	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.
<code>pf.console.title</code>	Defines the browser window or tab title for the administrative console, used to make separate instances identifiable.
<code>pf.console.session.timeout</code>	Defines the length of time in minutes until an inactive administrative console times out. The minimum setting is 1 minute; maximum is 8 hours (480 minutes). Default is 30 minutes.
<code>pf.log.eventdetail</code>	Enables or disables (the default) detailed event logging for actions performed by administrative-console users.
<code>pf.console.login.mode</code>	Indicates whether more than one administrative user may access the administrative console at one time. Supported values: Single Multiple. The default value is Multiple.
<code>pf.console.authentication</code>	Indicates whether administrators log on to PingFederate using credentials managed internally by PingFederate or externally by other systems.
<code>pf.admin.api.authentication</code>	Defines the authentication method of the PingFederate administrative API.
<code>ldap.properties.file</code>	When LDAP administrative-console authentication is enabled, indicates the name of the file containing configuration properties.
<code>cert.properties.file</code>	When certificate-based console authentication is enabled, indicates the name of the file containing configuration properties.
<code>radius.properties.file</code>	When RADIUS-based console authentication is enabled, indicates the name of the file containing configuration properties.
<code>pf.http.port</code>	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.
<code>pf.https.port</code>	Defines the port on which PingFederate listens for encrypted HTTPS (SSL/TLS) traffic. The default value is 9031.
<code>pf.secondary.https.port</code>	Defines a secondary HTTPS port that can be used, for example, with SOAP or artifact SAML bindings or for WS-Trust STS calls. To use this port, change the placeholder value to the port number you want to use.  Important: If you are using mutual SSL/TLS for either WS-Trust STS authentication or for SAML back-channel authentication, you <i>must use</i> this

Property	Description
	<p>port for security reasons (or use a similarly configured new listener) with either the <code>WantClientAuth</code> or <code>NeedClientAuth</code> parameter set to <code>true</code> in the <code>jetty-runtime.xml</code> file.</p> <p>For more information, see the note at the end of this table.</p>
<code>pf.engine.bind.address</code>	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.
<code>pf.monitor.bind.address</code>	Defines the IP address over which an SNMP agent and JMX communicate with PingFederate. Use for deployments where multiple network interfaces are installed on the machine running PingFederate.
<code>pf.engine.prefer_ipv4</code>	Defines the protocol to be used by PingFederate. <code>True</code> (the default) enables use of IPv4 only. <code>False</code> enables use of both IPv4 and IPv6.
<code>http.proxyHost</code> and <code>http.proxyPort</code>	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.
<code>https.proxyHost</code> and <code>https.proxyPort</code>	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.
<code>http.nonProxyHosts</code>	<p>Specifies one or more destinations where PingFederate is not required to proxy its HTTP <i>and</i> HTTPS traffic through the forward proxy server configure by the <code>http[s].proxyHost</code> and <code>http[s].proxyPort</code> properties. This property supports multiple values separated by the pipe character () and the wildcard character (*) for pattern matching; for example:</p> <p><code>*.example.com localhost</code></p>
<code>pf.runtime.context.path</code>	<p>Allows customization of the server path for PingFederate endpoints.</p> <p> Note: If this property is changed, the path must also be added to the base URL for your server (as defined on the Server Configuration > Server Settings > Federation Info screen).</p>
<code>pf.log.dir</code>	<p>Network path to the output location of log files. The default is:</p> <p><code><pf_install>/pingfederate/log</code></p>
<code>pf.hsm.mode</code>	Enables or disables (the default) a FIPS-compliance Hardware Security Module (HSM).
<code>pf.hsm.hybrid</code>	<p>Enables or disables the HSM hybrid mode. Applicable only when the <code>pf.hsm.mode</code> property is configured to use an HSM.</p> <p>When set to <code>true</code>, keys and certificates can be stored on either the HSM or the local trust store. When set to <code>false</code> (the default), keys and certificates on are stored on the HSM when applicable.</p> <p>The HSM hybrid mode allows an organization 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 Transition to an HSM on page 193.</p>
<code>pf.provisioner.mode</code>	Enables or disables (the default) outbound provisioning. Also used to enable provisioning failover.
<code>pf.heartbeat.system.monitoring</code>	Enables or disables (the default) the heartbeat endpoint (<code>/pf/heartbeat.ping</code>) to return detailed system monitoring information through a customizable Velocity

Property	Description
	template file (<pf_install>/pingfederate/server/default/conf/template/heartbeat.page.template).
	Note: Additional configuration of the listener ports (including adding new listeners) is available via 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 preconfigured SSL secondary port, the WantClientAuth parameter is set to true and the NeedClientAuth parameter is set to false by default.)
	1. Edit the <pf_install>/pingfederate/bin/run.properties file.
	Consider creating a backup copy of the file.
	2. Modify the applicable properties.
	3. Restart PingFederate.
	Important: For a clustered PingFederate environment, changes made runtime-server properties must be applied to other server nodes manually. The configuration replication process does not push settings in the run.properties file to other server nodes in the cluster.

PingFederate log files

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 purpose.

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 server to use the same logging configuration, manually synchronize the `log4j2.xml` file across multiple PingFederate server.

Log levels and verbosity

Log messages are categorized into six log levels:

1. FATAL
2. ERROR
3. WARN
4. INFO
5. DEBUG
6. TRACE

Starting with version 8.2, PingFederate only records messages that are tagged with log level INFO, WARN, ERROR, and FATAL to the server log (and the provisioner log). Messages that are tagged DEBUG (or TRACE) are not recorded to optimize performance. Console logging is also disabled for the same reason.

For troubleshooting purpose, you may adjust the log level to DEBUG in the `log4j2.xml` file and optionally re-enable console logging.



Important: When debug messages and console logging are no longer required, ensure they are turned 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 [Enable debug messages and console logging](#) on page 513.

Updating the log level for existing `Logger` elements does not require a restart. The adjustments are activated within half a minute. Other changes require a restart of PingFederate; for example:

- Adding new `logger` elements
- Updating the `size` attribute in the `RollingFile/Policies/SizeBasedTriggeringPolicy` element
- Updating the `max` attribute in the `RollingFile/DefaultRolloverStrategy` element
- Updating the `filePattern` attribute without modifying the `fileName` attribute in the `RollingFile` appender

Fields (and attributes)

Some logs, such as the audit log and the administrative API log, can be customized to log additional (or less) information by modifying their `pattern` elements. Available fields are documented inline in the `log4j2.xml` file.



Tip: PingFederate can be configured to log user attributes (if they are present) in the audit log, transaction log, and server log. When privacy is required 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 useful for troubleshooting and support purposes to aggregate and analyze log entries tied to the same original request. The tracking ID (`%X{trackingid}`) may also be added to the configuration for the transaction log, or be 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 CEF format. Furthermore, the audit log may also be written in a format that can be 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-event-detail.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 may 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. As needed, administrators may adjust the file size and the number of files to be retained in the `log4j2.xml` configuration file.

For more information about Log4j 2, please refer to the [Log4j 2 open-source project](http://logging.apache.org/log4j/2.x/manual/index.html) (`logging.apache.org/log4j/2.x/manual/index.html`).

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)

These configuration files are located in the `<pf_install>/pingfederate/etc` directory and independently managed 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, Log4j 2 configuration settings (in the `<pf_install>/pingfederate/server/default/conf/log4j2.xml` file) may be adjusted to deliver the desired level of detail surrounding each event.

Events logged by PingFederate includes (but 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

- IdP/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

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 the corresponding **Save** button in the administrative console is clicked. 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 role(s) assigned to the administrator at the time the event occurred.
- The type of event that occurred.
- Basic information about the event.

Each of the above fields is separated by a vertical pipe (|) for easier parsing.

Detailed event logging

PingFederate® can also be configured to log additional event information to a separate log file. When detailed event logging is enabled, besides writing basic information to `<pf_install>/pingfederate/log/admin.log`, PingFederate logs detailed information about each event to `admin-event-detail.log` (in the same log directory). Events between `admin.log` and `admin-event-detail.log` are linked by a unique event ID. Each entry in `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.

 **Note:** 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). Actions performed through these endpoints are logged for auditing purposes, as described below:

API	Port	Log File
Administrative API	Administrative Port	<code>admin-api.log</code>
OAuth Client Management Service	Runtime Port	<code>runtime-api.log</code>
OAuth Access Grant Management Service	Runtime Port	<code>runtime-api.log</code>
Session Revocation API	Runtime Port	<code>runtime-api.log</code>

Administrative API audit log

PingFederate® records actions performed via the administrative API. This information is recorded in the `<pf_install>/pingfederate/log/admin-api.log` file. While the events themselves are not

configurable, Log4j 2 configuration settings (in the `<pf_install>/pingfederate/server/default/conf/log4j2.xml` file) may 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 the above 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 `<pf_install>/pingfederate/log/runtime-api.log` file. While the events themselves are not configurable, Log4j 2 configuration settings (in the `<pf_install>/pingfederate/server/default/conf/log4j2.xml` file) may 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 the above 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 (inbound and outbound messages). Administrators may configure transaction logging to any of the four modes on a per-connection basis or override the logging mode for all SP connection, 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: <ul style="list-style-type: none"> • 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)

Mode	Description
	<ul style="list-style-type: none"> Target URL (if SSO transaction)
Enhanced	<p>Includes everything logged at the Standard level plus:</p> <ul style="list-style-type: none"> SAML_SUBJECT* Binding Relay state (if available) Signature policy Signature status HTTP request parameters (outbound messages only) <p>* Only when available in a SAML assertion, a single-logout request, an STS Request Security Token Response (RSTR), or an authentication request (AuthnRequest)</p>
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.

- To configure transaction logging mode on a per connection basis:
 - Select the applicable connection from the **IdP Configuration** or **SP Configuration** screen.
 - Click **General Info** and then select the one of the four logging modes.
- To override transaction logging mode for all SP (or IdP) connections:
 - On the **IdP Configuration** (or **SP Configuration**) screen, click **Manage All** under **SP Connections** (or **IdP Connections**).
 - Turn on the **Logging Mode Override** setting and select a logging mode for all SP (or IdP) 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. PingFederate records this information in the `audit.log` file, located in the `<pf_install>/pingfederate/log` directory. The following table describes the default elements that PingFederate writes to the audit log (in the order that they are listed):



Element	Description
%d	Transaction time.
trackingid	The tracking ID values uniquely identify user sessions, useful for correlating log messages in the audit and server logs.
event	<p>The type of transaction; for example, SSO, OAuth, AUTHN_ATTEMPT, and AUTHN_REQUEST.</p> <p>AUTHN_ATTEMPT and AUTHN_REQUEST indicate an authentication attempt against an IdP adapter instance and an authentication request sent to another IdP partner (through an IdP connection), respectively.</p>
subject	The subject of the transaction or authentication attempt.
ip	Incoming IP address.
app	The target SP application (when available).
connectionid	The connection identifier associated with the transaction.
protocol	The associated identity protocol; for example, SAML20 or OAuth20.
host	PingFederate host name or IP address.
role	The role of PingFederate played for the transaction.

Element	Description
status	The status of the transactions.
adapterid	The ID of an adapter instance. Consider adding the authenticationsourceid and targetsessionid elements to log additional information about the request.
description	Description of an authentication failure (when information is available from an IdP adapter).
responsetime	Time elapsed (in milliseconds) from when a final request for a transaction is received to when the audit message is written. This value serves as an approximation of total transaction processing time and may be useful for monitoring trends.

Each element is separated by a vertical pipe (|). Elements are configurable by editing the <pf_install>/pingfederate/server/default/conf/log4j2.xml file.

The following table describes other elements (in alphabetical order) that can be added to the audit log:

Element	Description
accessgrantguid	The GUID of the OAuth access grant (for OAuth transactions).
assertionid	The unique ID for the SAML assertion.
atrackingid	The tracking ID for OAuth access token. It could be used to analyze the flow of OAuth access tokens in the audit log and between PingFederate and PingAccess.
attributes	User attributes received (for an SP log) or sent (for an IdP log).
authenticationsourceid	An array of one or more IdP adapters or IdP connections (or both) invoked in an authentication or logout flow; for example, [idpAdapterOne, idpConnectionX]
granttype	OAuth grant type.
initiator	(SAML 2.0 only) The federation role that initiated the SSO or SLO: SP or IDP.
inmessagetype	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.
localuserid	The local ID used for the transaction (when account linking is enabled at the SP).
outxmlmsg	The outgoing message; for example, a SAML Response or the information pertaining to a response for an OAuth request.
pfversion	The PingFederate version.
requestid	The ID of a SAML request.
responseid	The ID of a SAML response.
requeststarttime	The start time of the request in milliseconds since midnight, January 1, 1970 UTC.
stspluginid	For WS-Trust STS transactions, the ID for the token-processor or token-generator instance.
targetsessionid	An array of one or more SP adapters or SP connections (or both) invoked in an authentication or logout flow.
validatorid	The ID of the Password Credential Validator instance (for the successful attempts).



Element	Description
virtualserverid	The virtual server ID of a request (if applicable).
	Tip: Elements from the audit log may be output to different formats, including databases, CEF, and Splunk. For information, see Logging in other formats on page 104 .
	Tip: The audit log records SSO, SLO, OAuth AS, WS-Trust STS, and SCIM inbound provisioning transactions. Outbound provisioning transactions are not included; they are logged to the outbound provisioning audit log; the <code>provisioner-audit.log</code> file in the <code><pf_install>/pingfederate/log</code> directory.

Outbound provisioning audit logging

The PingFederate® `provisioner-audit.log` records outbound provisioning transactions, intended to facilitate security auditing. Elements of this log are described in the following table:



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 <code>true</code> ; otherwise, the value is <code>false</code> .
non_success_cause	Description of failure cause.

Configuration is maintained in the `log4j2.xml` file located in the `<pf_install>/pingfederate/server/default/conf` directory.

-  **Tip:** Elements from the provisioner audit log may be output to different formats, including database and Splunk. For information, see [Logging in other formats](#) on page 104.
-  **Important:** Monitoring outbound provisioning transactions using JMX has been deprecated. PingFederate logs outbound provisioning events to the `provisioner-audit.log` log file out-of-the-box.

Server logging

When PingFederate is configured to log `DEBUG` messages (for troubleshooting purpose), it records all runtime and administrative events, including status and error messages that can be used for troubleshooting, in the `<pf_install>/pingfederate/log/server.log` file.

-  **Note:** `DEBUG` messages are turned off by default. For more information, see [Enable debug messages and console logging](#) on page 513.
-  **Tip:** Administrators may choose write server log messages to a database server (see [Logging in other formats](#) on page 104).

Elements of this log are described in the following table:

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.

Item	Description
%p	Logging level.
%c	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 that administrators can use to filter server logs: `logfilter.bat` or `logfilter.sh`, located in the `<pf_install>/pingfederate/bin` directory.

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
<code>-entityid <entity ID></code>	These two commands must be used together and return a list of transactions for the specified federation partner's entity ID and transaction subject.
<code>-subject <subject></code>	
<code>-trackingid <tracking ID></code>	This command returns a list of transactions with the same tracking ID.
<code>-sessionxrefid <session cross-reference ID></code>	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 <code><session cross-reference ID></code> may be the value of any of the following transaction tags in the target server log(s): <ul style="list-style-type: none"> • Artifact • Session Index • Assertion ID

Server log filter parameters (optional)

Parameter	Description
<code>-logsdir <log files directory></code>	Full or relative path to source directory for the logs. Default: all <code>server.log</code> files are written to the <code><pf_install>/pingfederate/log</code> directory (a setting that can be adjusted by the <code>pf.log.dir</code> property in the <code><pf_install>/pingfederate/bin/run.properties</code> file.
<code>-outputfile <output file></code>	Output path and file for the returned list. Default: <code>\$pf.log.dir/logfilter_output.log</code> .
<code>-outputtoconsole</code>	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. Administrators can control settings for this log, as needed, in the file `logfilter.log4j2.xml`, located in the `<pf_install>/pingfederate/bin` directory.

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 may 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.

Write 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. These scripts are located in the `<pf_install>/pingfederate/server/default/conf/log4j/sql-scripts` directory.



Note: PingFederate has been tested with vendor-specific JDBC 4.1 drivers. To obtain your database driver JAR file, contact your database vendor. Database driver file should be installed to the `<pf_install>/pingfederate/server/default/lib` directory. You must restart the server after installing the driver.

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.

1. Edit `<pf_install>/pingfederate/server/default/conf/log4j2.xml`.
2. Under the **Preserve messages in a local file** section, for each log that you want to enable database logging, uncomment the preset JDBC appender configuration based on the choice of your database server.

Audit log

- `SecurityAuditToMySQLDB` (for Oracle MySQL)
- `SecurityAuditToOracleDB` (for Oracle Database)
- `SecurityAuditToPostgreSQLDB` (for PostgreSQL)
- `SecurityAuditToSQLServerDB` (for Microsoft SQL Server)

Provisioner audit log

- `OutboundProvisionerEventToMySQLDB` (for Oracle MySQL)
- `OutboundProvisionerEventToOracleDB` (for Oracle Database)
- `OutboundProvisionerEventToPostgreSQLDB` (for PostgreSQL)
- `OutboundProvisionerEventToSQLServerDB` (for Microsoft SQL Server)

Provisioner log


- `ProvisionerLogToMySQLDB` (for Oracle MySQL)
- `ProvisionerLogToOracleDB` (for Oracle Database)
- `ProvisionerLogToPostgreSQLDB` (for PostgreSQL)
- `ProvisionerLogToSQLServerDB` (for Microsoft SQL Server)

Server log

- `ServerLogToMySQLDB` (for Oracle MySQL)
- `ServerLogToOracleDB` (for Oracle Database)
- `ServerLogToPostgreSQLDB` (for PostgreSQL)
- `ServerLogToSQLServerDB` (for Microsoft SQL Server)




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 reasons. Both appenders must also be enabled (uncommented).

 **Tip:** Review inline comments and notes in the `log4j2.xml` file for more information about each appender.

3. Replace placeholder parameter values in `log4j2.db.properties` in the same `conf` directory for the applicable 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 either `obfuscate.sh` or `obfuscate.bat`, located in `<pf_install>/pingfederate/bin`. 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:

- `org.sourceid.websso.profiles.sp.SpAuditLogger` (Browser SSO SP and adapter-to-adapter)
- `org.sourceid.websso.profiles.idp.IdpAuditLogger` (Browser SSO IdP and adapter-to-adapter)
- `org.sourceid.websso.profiles.idp.AsAuditLogger` (OAuth authorization server)
- `org.sourceid.wstrust.log.STSAuditLogger` (WS-Trust STS, IdP and/or SP)

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 log

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


The Common Event Format (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 via HP ArcSight tools. Alternatively, administrators can write the information to a flat file in CEF; however, using syslog, when available, is recommended.


 **Note:** PingFederate is certified with HP ArcSight for interoperability using the default elements defined in `log4j2.xml`. Any additions to these elements may render your CEF logging incompatible with HP ArcSight.

Write audit log in CEF


1. Edit `<pf_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)

 **Note:** The `SecurityAuditToCEFSyslog` `Socket` appender is followed by two related appenders: `PingFailover` and `RollingFile`. Together, they create a running `audit-cef-syslog-failover.log` file in the log directory in the event that CEF logging fails for any reason. Both appenders must also be enabled (uncommented).

 **Tip:** 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:
 - `org.sourceid.websso.profiles.sp.SpAuditLogger` (Browser SSO SP and adapter-to-adapter)
 - `org.sourceid.websso.profiles.idp.IdpAuditLogger` (Browser SSO IdP and adapter-to-adapter)
 - `org.sourceid.websso.profiles.idp.AsAuditLogger` (OAuth authorization server)
 - `org.sourceid.wstrust.log.STSAuditLogger` (WS-Trust STS, IdP and/or SP)


 **Important:** As indicated in the **IMPORTANT** comments for the loggers, you must also remove some of the existing appender references.

Write provisioner audit log in CEF


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)
- `OutboundProvisionerEventToCEFFile` (a `RollingFile` appender under the **Outbound provisioner audit log for CEFFile** section)

 **Note:** The `OutboundProvisionerEventToCEFSyslog` `Socket` appender is followed by two related appenders: `PingFailover` and `RollingFile`. Together, they create a running `provisioner-audit-cef-syslog-failover.log` file in the log directory in the event that CEF logging fails for any reason. Both appenders must also be enabled (uncommented).

 **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.


Write audit log for Splunk

Splunk is enterprise software that allows for monitoring, reporting, and analyzing consolidated log files. Splunk captures and indexes real-time data into a single searchable repository from which reports, graphs, and other data visualization can be generated.

Ping Identity provides a custom Splunk App for PingFederate® to process audit logs generated by a PingFederate deployment. The application provides rich system monitoring and reporting, including:

- Current transaction and system reports
- Service reports such as a daily usage report and IdP and SP reports per connection
 -  **Note:** OAuth AS and WS-Trust STS transactions are not currently supported for Splunk.
- Trend reports such as weekly and monthly usage reports, and trend analysis

The application uses a specially formatted version of the audit log (`splunk-audit.log`), which is written to the PingFederate log directory when the setup steps described below are followed.

 **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. Download the free application from splunkbase.splunk.com. (Search for PingFederate.)

1. Set up your Splunk server.

- a) If you have not done so, download and install Splunk.
- b) Enable a receiver to *listen* for data from the PingFederate server.

For more information, please refer to [Splunk documentation](http://docs.splunk.com/Documentation/Forwarder/6.5.0/Forwarder/HowtoforwarddatatoSplunkEnterprise) (docs.splunk.com/Documentation/Forwarder/6.5.0/Forwarder/HowtoforwarddatatoSplunkEnterprise).

- c) 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) Look for the following Logger elements located under the **Loggers** section:

- `org.sourceid.websso.profiles.sp.SpAuditLogger` (Browser SSO SP and adapter-to-adapter)
- `org.sourceid.websso.profiles.idp.IdpAuditLogger` (Browser SSO IdP and adapter-to-adapter)
- `org.sourceid.websso.profiles.idp.AsAuditLogger` (OAuth authorization server)
- `org.sourceid.wstrust.log.STSAuditLogger` (WS-Trust STS, IdP and/or SP)


- c) Uncomment the SecurityAudit2Splunk RollingFile appender reference (`<appender-ref ref="SecurityAudit2Splunk"/>`) from the one or more of the Logger elements.

For example, the default logger for an IdP audit log reads:

```
<Logger name="org.sourceid.websso.profiles.idp.IdpAuditLogger"
  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"
  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.

3. Set up Splunk Universal Forwarder.

- a) Download the Splunk Universal Forwarder from [Splunk](http://www.splunk.com/en_us/download/universal-forwarder.html) (www.splunk.com/en_us/download/universal-forwarder.html) 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, please refer to [Splunk documentation](http://docs.splunk.com/Documentation/Forwarder/6.5.0/Forwarder/HowtoforwarddatatoSplunkEnterprise) (docs.splunk.com/Documentation/Forwarder/6.5.0/Forwarder/HowtoforwarddatatoSplunkEnterprise).


Account management

The PingFederate® administrative console supports four authentication schemes:

- Native authentication
- LDAP authentication
- RADIUS authentication
- Certificate-based authentication

For native authentication, the administrative console offers two system administration styles:

- **Single-User Administration** restricts access to only one account with full privileges.
- **Multi-User Administration** allows the creation of multiple console accounts with different roles. This is the default style.


 **Tip:** You can update administration style in the **Server Configuration > Server Settings > System Administration** screen.

For role-based access control, PingFederate provides two account types and three administrative roles, as shown in the following table:

PingFederate User Access Control

Account type	Administrative role	Access privileges
Admin	Admin	Configure partner connections and most system settings (except the management of native accounts and the handling of local keys and certificates).
Admin	Crypto Admin	Manage local keys and certificates.

Account type	Administrative role	Access privileges
Admin	User Admin	Create users, deactivate users, change or reset passwords, and install replacement license keys.
Auditor	<i>Not applicable</i>	View-only permissions for all administrative functions. When the Auditor role is assigned, no other administrative roles may be set.


 **Note:** All three administrative roles are required to access and make changes through the following services:

- The `/ConfigArchive` administrative API endpoint
- The **Server Configuration** > **Configuration Archive** screen in the administrative console
- The `/pf-mgmt-ws/ws/ConnectionMigrationMgr` and `/pf-mgmt-ws/ws/ConfigReplication` Connection Management Web Service endpoint
- The **Server Configuration** > **Application Authentication** screen in the administrative console for the **Connection Management** service


For native authentication, access and authorization are controlled by the local accounts defined in the **Server Configuration** > **Account Management** screen (unless the administration style has been updated to the **Single-User Administration**).

As needed, you can switch from native authentication to an alternative console authentication. Note that access and authorization are defined in the respective configuration file.

An administrative user may log on from more than one browser or location. Moreover, multiple administrative users can log 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 `<pf_install>/pingfederate/bin/run.properties` file. Regardless of the property configuration, any number of auditors may log on at any time.

 **Note:** For security, after three failed logon attempts from the same location within a short time period, the administrative console will temporarily lock out further attempts by the same user. The user must wait one minute to try again.

Note that the accounts in the **Account Management** screen 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 **Account Management** screen appears only if the administrative API is left to use native authentication, and vice versa.

 **Tip:** If you have connected PingFederate to PingOne, you may also single sign-on from the PingOne admin portal to the administrative console.

Enable native authentication

When the administrative console is protected by native authentication, access is restricted to the local accounts defined in the **Server Configuration** > **Account Management** screen.

1. In the `<pf_install>/pingfederate/bin/run.properties` file, change the value of the `pf.console.authentication` property as shown below:






```
pf.console.authentication=native
```

2. Start or restart PingFederate®.

Manage native accounts and role assignments

1. Go to the **Server Configuration** > **Account Management** screen.


Task	Steps
Create a native account	<ol style="list-style-type: none"> 1. On the Account Management screen, click Create User. 2. On the User Information screen, enter a username and other optional information.

Task	Steps
	<p> Note: If you want PingFederate® to notify the user about password changes via email, you must supply an email address.</p> <p>3. On the Password Generation screen, enter a password or click Generate one-time password to generate a random password for the account.</p> <p> Note: Upon successful authentication, the user will be required to change the password of the account immediately.</p> <p>4. On the Summary screen, review your configuration, modify as needed, and then click Done.</p> <p>5. On the Account Management screen, select the applicable account type (Auditor or Admin) and one or more administrative roles for an Admin account.</p> <p>6. Repeat these steps to create additional accounts.</p>
Modify user information	<p>1. On the Account Management screen, select the account by its username.</p> <p> Note: Applicable only to active accounts.</p> <p>2. On the User Information screen, update the record, and then click Done.</p> <p>3. Repeat these steps to update other accounts.</p>
Update role assignments	<p>1. Select a different account type (Auditor or Admin) for one or more accounts.</p> <p>2. Select or clear the check boxes that correspond to the three administrative roles (User Admin, Admin, and Crypto Admin for one or more accounts.</p> <p> Note: Applicable only to the Admin accounts.</p>
Deactivate or reactive a native	<p>1. Click Deactivate or Activate under Action.</p> <p>2. Repeat this step to deactivate or reactive other accounts.</p> <p> Note: For traceability and accountability purposes, native accounts cannot be deleted; their records are retained and they can be reactivated if needed.</p>


2. Click **Save** to retain your configuration.

Enable notification for account management events

Administrators may enable email notification for account management events.

 **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.

 **Note:** If you are using an alternative console authentication for the administrative console or the administrative API, password notifications (if any) is handled by the third-party system.

1. Go to the **Server Configuration > Account Management** screen.
2. Select the **Notify Administrator of Account Change** check box.

An email address must be provided for the applicable accounts.
3. If you have not yet configured PingFederate® to use your email server, click **Email Server Settings** and complete the configuration.
4. Click **Save** to retain your configuration.

When enabled, email notification is triggered by the following events:

Event	Notification
An administrator has turned off the Notify Administrator of Account Changes option.	An email notification is sent immediately to all administrators. The message includes the username of the administrator who made the change.
An administrator's email address has been updated by another administrator.	An email notification is sent immediately to the previous email address and the updated email address. The message includes the username of the administrator who made the change.
An administrator's password has been changed.	An email notification is sent immediately to the administrator whose password has been changed. The message includes the username of the administrator who made the change.

Administrators should review the changes and ensure that the changes legitimate.

Set or reset password

User administrators generate temporary passwords as they create new native accounts for new users in the **Password Generation** screen.

User administrators can also reset and assign temporary passwords for existing users who forget their passwords. Upon successful authentication, the users are required to change their passwords immediately.



Note: If you are using an alternative console authentication for the administrative console or the administrative API, password management is handled by the third-party system.

1. Go to the **Server Configuration > Account Management** screen.
2. Optional: Select the **Notify User of Password Change** check box if you want PingFederate® to notify the user about password changes via email.
3. Click **Reset Password** under **Action** for the applicable account.
4. On the **Password Generation** screen, enter a password or click **Generate one-time password** to generate a random password for the account, and then click **Save**.

After you click **Save**, if you have enabled email notification for password changes, the new password is emailed to the user automatically.

Change password

Administrative users and auditors can change the passwords of their native accounts at any time.



Note: 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.

1. Go to the **Server Configuration > Account Management** screen.
2. Click **Change Password** under **Action**.
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. Click **Save** to retain your configuration.

Alternative console authentication

As an alternative to using PingFederate's own internal data store for authentication to the administrative console, you can configure PingFederate® to use either your network's LDAP user-data store, the RADIUS protocol, or client certificates. You can configure any of these alternatives at any time.

Note that 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 the **Server Configuration > Account Management** screen, 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	Admin	Configure partner connections and most system settings (except the management of native accounts and the handling of local keys and certificates).
Admin	Crypto Admin	Manage local keys and certificates.
Admin	User Admin	Create users, deactivate users, change or reset passwords, and install replacement license keys.
Auditor	<i>Not applicable</i>	View-only permissions for all administrative functions. When the Auditor role is assigned, no other administrative roles may be set.



Note: All three administrative roles are required to access and make changes through the following services:

- The `/ConfigArchive` administrative API endpoint
- The **Server Configuration > Configuration Archive** screen in the administrative console
- The `/pf-mgmt-ws/ws/ConnectionMigrationMgr` and `/pf-mgmt-ws/ws/ConfigReplication` Connection Management Web Service endpoint
- The **Server Configuration > Application Authentication** screen in the administrative console for the **Connection Management** service

Enable LDAP authentication

The LDAP authentication setup is available via configuration files in the `<pf_install>/pingfederate/bin` directory.



Note: When LDAP authentication is configured, PingFederate® does not lock out administrative users based upon the number of failed logon attempts. Responsibility for preventing access is instead delegated to the LDAP server and enforced according to its password lockout settings.

1. In the `<pf_install>/pingfederate/bin/run.properties` file, change the value of the `pf.console.authentication` property as shown below:


```
pf.console.authentication=LDAP
```
2. In the `<pf_install>/pingfederate/bin/ldap.properties` file, change property values as needed for your network configuration.

See the comments in the file for instructions and additional information.

Note that the roles configured in the properties file apply to both the administrative console and the administrative API.




Important: Be sure 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.


-  **Tip:** You can also use this configuration file in conjunction with RADIUS authentication to determine permissions dynamically via an LDAP connection.

3. Start or restart PingFederate.

Enable RADIUS authentication

The RADIUS authentication setup is available via configuration files in the `<pf_install>/pingfederate/bin` directory. The RADIUS protocol provides a common approach for implementing strong authentication in a client-server configuration. PingFederate® supports the protocol scenarios for one-step authentication (by appending to the password a one-time passcode obtained from an authenticator, for instance) and two-step authentication (through a challenge-response process, for example).

 **Note:** When RADIUS authentication is configured, PingFederate does not lock out administrative users based upon the number of failed logon attempts. Responsibility for preventing access is instead 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.

1. In the `<pf_install>/pingfederate/bin/run.properties` file, change the value of the `pf.console.authentication` property as shown below:

```
pf.console.authentication=RADIUS
```
2. In the `<pf_install>/pingfederate/bin/radius.properties` file, change property values as needed for your network configuration.

See the comments in the file for instructions and additional information.

Note that 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 (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 (also in the `bin` directory) to map LDAP group-based permissions to PingFederate roles.

3. Start or restart PingFederate.

Enable certificate-based authentication


To enable client-certificate authentication, PingFederate® administrative users must have imported into their web browsers an X.509 key and certificate suitable for user authentication. In addition, the corresponding root CA certificate(s) must be contained in the Java runtime or the PingFederate trusted store.

Other setup steps, including designating user permissions, are required via configuration files in the `<pf_install>/pingfederate/bin` directory.

Note that the roles configured in the properties file apply to both the administrative console and the administrative API.

1. If not already done, import the necessary client key and certificate into the web browser used to access PingFederate.
2. Log on normally (using username and password) to the PingFederate console as a user with permissions that include the **Crypto Admin** role.
3. 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).

You can import a certificate to PingFederate in the **Server Configuration > Trusted CAs** screen.

 **Tip:** You may wish to click the Serial number and copy the Issuer DN to use in a couple steps later.

4. In the `pingfederate/bin/run.properties` file, change the value of the `pf.console.authentication` property as shown below:

```
pf.console.authentication=cert
```

5. 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.

If you copied the Issuer DN a couple steps earlier, paste this value.


See the comments in the file for instructions and additional information.

Note that the roles configured in the properties file apply to both the administrative console and the administrative API.

6. Repeat the previous step for any additional CAs as needed.
7. 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.


8. Repeat the previous step for all users as needed.

 **Note:** Other settings in the properties file are used to display the user's ID (the Subject DN) in abbreviated form in the administrative console.

9. Start or restart PingFederate.

Manage email configuration

If you are using email notifications for password resets, licensing events, certificate-expiration warnings, or automatic metadata-reloading events, you must set up and maintain a connection to the email server that PingFederate® uses to send messages.

 **Tip:** In a clustered PingFederate environment, the email notifications are only sent by one of the nodes in the cluster: the console node or any engine node. If this node leaves the cluster (planned, or not), another node picks up this task automatically.


1. Go to the **Server Configuration > Email Configuration** screen.

Applicable only if you have already configured email notifications in the **Server Configuration > Server Settings > Runtime Notifications** or **Account Management** screen.

When you configure email notifications for the first time in the **Runtime Notifications** or **Account Management** screen, the administrative console guides you to the **Email Configuration** screen to set up an email server.

2. Configure the email server settings as follows:

Field	Description
“From” Address (Required)	The email address that appears in the “From” header line in email messages generated by PingFederate. The address must be in valid format but need not be set up on your system.
Email Server (Required)	The IP address or hostname of your email server.
SMTP Port (Required)	The SMTP port on your email server. The default value is 25.
Encryption Method (Required)	Select SSL/TLS to establish a secure connection to the email server at the SMTPS port.

Field	Description
	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) to establish an unencrypted connection to the email server at the SMTP port.
SMTPS Port	The secure SMTP port on your email server. This field is not active unless SSL/TLS is the chosen encryption method. The default value is 465.
Verify Hostname	Indicates whether to verify the hostname of the email server matches the Subject (CN) or one of the Subject Alternative Names from the certificate. Visible and selected when either SSL/TLS or STARTTLS is the chosen encryption method.
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.
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 may contain sensitive information, to the server log. Consider enabling debug messages solely for troubleshooting purpose and disabling this option when debug messages are no longer required.
Username	Authorized email account.
Password	Password for the authorized email account.
Test Address	Enter an email address the PingFederate should use to verify connectivity with the configured email server.

- Optional: Click **Test Email Connectivity** to verify connectivity with the configured server.

Although optional, it is recommend that you run a connectivity test. A message next to the button indicates a successful test. Verify that the test email address received a message from the server. Test reports are written to the server log, located in the `<pf_install>/pingfederate/log` directory.

- Click **Save** to commit the email server configuration.

Configure virtual host names

In certain contexts, the SAML specifications require that XML messages include a URL identifying the host name to which the sender directed the message. (The name of the XML element containing the URL varies among protocols.) In addition, the recipient must verify that the value matches the location where the message is received.

Depending on your networking requirements, this specification can present problems; for example, in the case of proxy forwarding, where the final destination host name might be unknown to your federation partner. To provide more flexibility in such cases, you can set up a list of alternative host names on the **Server Configuration > Virtual Host Names** screen for PingFederate® to use as part of its message-security validation.


Note that virtual host names are used for a different purpose than virtual server IDs, which provide separate unique identifiers on a per-connection basis for a federation deployment, normally in the *same* domain. Depending on your needs, however, you can configure virtual server IDs and virtual hosts in the same installation of PingFederate.

- To add a new entry, enter the desired value and click **Add**.
- To modify an existing entry, use the **Edit**, **Update**, and **Cancel** workflow.
- To remove an existing entry, use the **Delete** and **Undelete** workflow.

Export metadata to an XML file

For SAML deployments, if your partners do not support consuming metadata by URL, you can export metadata for any connection or select certain non connection-specific metadata for export. Although optional, it is recommended to sign the metadata, so that partners can verify the authenticity of the metadata.

1. Go to the **Server Configuration > Metadata Export** screen.
2. On the **Metadata Role** screen, select the applicable role.
3. On the **Metadata Mode** screen, select the information to be included in the metadata and whether to use the SOAP channel.
 - a) Select the **Use a connection for metadata generation** option if you want to export metadata for a specific connection; otherwise select the **Select information to include in metadata manually** option to create a metadata export file that is not bounded to a specific connection.
 - b) Select the **Use the secondary port for SOAP channel** check box if the PingFederate® secondary HTTPS port is configured and you want to use it for the SOAP channel.


 **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. Follow the rest of the workflow to export a metadata XML file, including selecting the certificate to sign the metadata XML file in the **Metadata Signing** screen (as needed).
5. Pass the metadata XML file to one or more partners.


Provide SAML metadata by file

You can export metadata for any SAML connection to an XML file. This is useful in a situation, where you have already created a SAML connection to the partner and your partner prefers consuming SAML metadata by file.

1. Go to the **Server Configuration > Metadata Export** screen.
2. On the **Metadata Role** screen, select the applicable role.
3. On the **Metadata Mode** screen, select the information to be included in the metadata and whether to use the SOAP channel.
 - a) Choose the **Use a connection for metadata generation** option.
 - b) Select the **Use the secondary port for SOAP channel** check box if the PingFederate® secondary HTTPS port is configured and you want to use it for the SOAP channel.

 **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** screen, select the applicable connection from the list.

 **Note:** 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 connection metadata are specific to the selected virtual server ID. Re-export the connection metadata for your partners after updating your virtual server IDs.

5. Optional: On the **Metadata Signing** screen, 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** workflow 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** screen, click **Export** to save the metadata XML file.
7. Pass the metadata XML file to your partner.

Provide general SAML metadata by file

You can manually select the desired information and generate a metadata XML file. This is useful for the following situations:

- You have not yet created a SAML 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 generate a non connection-specific metadata with selected information, which can be passed to multiple partners to expedite their configurations.

1. Go to the **Server Configuration > Metadata Export** screen.
2. On the **Metadata Role** screen, select the applicable role.
3. On the **Metadata Mode** screen, select the information to be included in the metadata and whether to use the SOAP channel.
 - a) Choose the **Select information to include in metadata manually** option.
 - b) Select the **Use the secondary port for SOAP channel** check box if the PingFederate® secondary HTTPS port is configured and you want to use it for the SOAP channel.



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** screen, select the applicable federation protocol that supports metadata exchange from the list.

If you have only enabled one protocol that supports metadata exchange in the **Server Configuration > Server Settings > Roles & Protocols** screen, this is a read-only screen with such protocol.

5. Optional: On the **Attribute Contract** screen, add one or more attributes.

Click **Edit**, **Update**, **Cancel**, or **Delete** under **Action** to modify or remove any attributes.

6. Optional: Select a signing key from the list.

When selected, the metadata export contains the public key that your partners can use to verify the digital signature of your SAML messages.

7. Optional: On the **Metadata Signing** screen, 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** workflow 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.

8. Optional: On the **XML Encryption Certificate** screen, select the certificate that your partner can use to encrypt XML content.

Applicable only to SAML 2.0.

If you have not yet created or imported your certificate into PingFederate, click **Manage Certificates** and use the **Certificate Management** workflow to complete the task.

9. On the **Export & Summary** screen, click **Export** to save the metadata XML file.
10. Pass the metadata XML file to one or more partners.

Sign XML files

PingFederate® supports digital signing of SAML metadata files that you and your partner might want to exchange. Applying a digital signature to an XML file ensures that the file is from the original source and that its contents have not been modified by a third party.

If you have previously exported an unsigned metadata XML file, you can create a (new) signed metadata XML file based on the unsigned metadata.

1. Go to the **Server Configuration > XML File Signatures** screen.
2. On the **Select XML File** screen, choose your file.
3. On the **Digital Signature Settings** screen, choose the certificate containing your signing key from the list.

 **Note:** By default, certificate and public-key information is included in the signed XML file. If you do not wish to include this information, clear the "... <KeyInfo> ..." and "... <KeyValue> ..." check boxes.

4. Select a signing algorithm corresponding to the selected certificate. Choices include RSA-SHA1, SHA256, SHA384, and SHA512; ECDSA-SHA256, SHA384 and SHA512; and SHA1.
5. On the **Export & Summary** screen, click the **Export** button to save the digitally signed file.

Replicate configuration

From the **Main Menu > Server Configuration > Cluster Management** screen, which is available only when the administrative console is running in cluster mode, you can replicate the current console configuration to all server nodes in the cluster.

This screen is also useful for verifying that nodes have joined the cluster.


For more information, refer to [Console configuration push](#) on page 551 in the PingFederate [Server clustering guide on page 533](#).

Configuration archive

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. These configuration archives can be used as backup files for the current PingFederate installation.

The automatic backup process typically completes without delays. For deployments with hundreds of connections or OAuth clients (or both), administrators can optionally configure PingFederate to create configuration archives periodically instead.

Additionally, administrators can export the current configuration to a ZIP file on the **Server Configuration > Configuration Archive** screen. This screen is only available to administrators, whose accounts have been assigned the User Admin, Admin, and Crypto Admin roles.

 **Caution:** As the backup file contains your complete PingFederate configuration, ensure the file is protected with appropriate security controls in place.

On the **Server Configuration > Configuration Archive** screen, administrators can also import an existing archive for immediate deployment into a running PingFederate server.


Furthermore, administrators can deploy a configuration archive manually by copying the ZIP file to `<pf_install>/pingfederate/server/default/data/drop-in-deployer` directory.


Configuration archives are intended for administrative-console configuration only. As such, the following files are *not* included in the archives:

- Launch scripts under the `<pf_install>/pingfederate/bin` and `<pf_install>/pingfederate/sbin` directories.
- Web container configuration files under the `<pf_install>/pingfederate/etc` directory.
- Log files under the `<pf_install>/pingfederate/log` directory.

- Database drivers and program files from adapters and any other plug-ins under 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) under the `<pf_install>/pingfederate/server/default/conf` directory.

If any changes have been made to files that are not part of the configuration archive, such files must be preserved manually.

 **Tip:** You may export a configuration archive, extract the ZIP file, and determine whether specific files are part of the configuration archive, or not.

 **Important:** Draft connections in archives are not imported. Complete any unfinished partner connections if you wish to include them in a full backup archive or in an archive to be used for configuration migration.

Configure a backup schedule


For deployments with hundreds of connections or OAuth clients (or both), administrators can optionally configure PingFederate® to create configuration archives periodically instead.

Edit the `org.sourceid.saml20.domain.mgmt.impl.DataArchiveBackup.xml` file, located in the `<pf_install>/pingfederate/server/default/data/config-store` directory.

Refer to the following table for each property.

Property	Description
ScheduledBackupEnabled	When set to <code>true</code> , PingFederate creates a configuration archive daily <i>if</i> configuration has changed since the creation of the last archive. The default value is <code>false</code> .
BackupTime	The local time at which PingFederate creates a configuration archive when the ScheduledBackupEnabled property is set to <code>true</code> . (Ignored when the ScheduledBackupEnabled is set to <code>false</code> .) The default value is <code>00:00:00</code> , 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 again immediately. The NumTriesWithIntegrityFailure property indicates the maximum number of attempts. (Ignored when the ScheduledBackupEnabled is set to <code>false</code> .) The default value is 3.
BackupOnAdminLogin	When set to <code>false</code> , PingFederate does not create a configuration archive when an administrator signs on. Note that regardless of the value of this property, PingFederate always create a configuration archive before an archive is imported. The default value is <code>true</code> .
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 instance, if your deployment has 50 connections, you can certainly 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` on the console node.

Export an archive

Administrators can export the current administrative-console configuration to a ZIP file on the **Server Configuration > Configuration Archive** screen. This screen is only available to administrators, whose accounts have been assigned the User Admin, Admin, and Crypto Admin roles.

On the **Export** screen, click **Export** and then save the ZIP file.



Caution: As the backup file contains your complete PingFederate configuration, ensure the file is protected with appropriate security controls in place.

Import an archive

Administrators can import a configuration archive (ZIP) on the **Server Configuration > Configuration Archive** screen. This screen 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 **Server Configuration > Configuration Archive > Import** screen, PingFederate® displays error messages if there are any missing plug-in components (such as adapters, database drivers, or token translators) on which the archive depends, or any mismatches of PingFederate licensing authorization. However, the administrator can choose to force the deployment and then install the necessary files later.



Note: 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.

1. On the **Import** screen, 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 plug-in components or other errors.

3. Click **Import**.

The administrative console will prompt you to confirm the import process.

To proceed, click **Yes**.

To abort, click **No** or **Cancel**.

Automating configuration migration

PingFederate® provides a configuration-migration tool that can be used for scripting the transfer of administrative-console configurations and configuration property files from one PingFederate server to another; for example, from a test environment to production. The tool may also be used to 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 Web Service in conjunction with an internal Web Service to export and import connections and other configurations, and to obtain lists (see [Connection Management Services](#)).



Important: The Connection Management Service must be activated for both the source and target servers before the `configcopy` tool can be used (see [Configure application authentication](#)).



Caution: For security reasons, the Connection Management Service should be disabled whenever it is not in use.

Copy 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. To do this, you copy the key (or keys, if you have more than one) 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 done before the first migration.

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.

Ensure you copy the entire key JSON message. For example, you would copy the text as shown in bold below:

```
{ "keys": [ { "kty": "oct", "kid": "j0PUEdAb95", "k": "AGi8Lg_ewdl-30Cx83kDMQE9oNlhgJSa_Pc4I8JTU8" } ] }
```

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:

```
{ "keys":
[ { "kty": "oct", "kid": "wER9zEpaPe", "k": "i0HQr9JmsqjAX4o_BQU1qGJzoLQI-nmwp8u3GyHzTB8" } ] }
```

Insert the comma and the source key as follows:

```
{ "keys":
[ { "kty": "oct", "kid": "wER9zEpaPe", "k": "i0HQr9JmsqjAX4o_BQU1qGJzoLQI-nmwp8u3GyHzTB8" },
{ "kty": "oct", "kid": "j0PUEdAb95", "k": "AGi8Lg_ewdl-30Cx83kDMQE9oNlhgJSa_Pc4I8JTU8" } ] }
```

Note that 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

For migrating data configured with the source server's administrative console, this tool performs these overall processing steps:

1. Retrieves specified connection, other configuration data (XML), or both, 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, we recommend using 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, outbound-provisioning channels, or data-store 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 data-store 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 may have been manually customized for the source configuration and need to be migrated. The file-copy command may also be used to copy the PingFederate internal, HyperSQL database when needed.

Managing Certificates

Administrators may use the `configcopy` tool to perform the following certificate-management tasks on the target PingFederate server:

- List source trusted 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

Migration tool

The `configcopy` tool may be used in conjunction with one or more property files to define the operational command and other parameters, including the source and/or target PingFederate servers, and to modify configuration settings as needed for the target environment.

Property-file templates are available for each command option. The template files are located in the `<pf_install>/pingfederate/bin/configcopy_templates` directory.



Note: Refer to the `README.txt` file in the `configcopy_templates` directory for a list of all commands and summary information. See the template files themselves for parameters associated with each command (or with use cases), as well as lists of Override Properties (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=<properties_file1>;
<properties_file2>;...
```

When paths are included with the file names, you cannot use backslashes (`\`). Use forward slashes (`/`) or escape the backslash (`\\`).

(On UNIX/Linux)

```
configcopy.sh -
Dconfigcopy.conf.file=<properties_file1>:<properties_file2>:...
```


Note that the file separators are platform specific, corresponding to the syntax used for system-level path separators.

Alternatively (or in addition), you can specify any property values via command-execution arguments, using the following syntax:


```
configcopy[.sh] -D<property>=<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 Web Services are password-protected. The usernames and passwords may be set in the properties file for both the source and target Web Services (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 always must be supplied where applicable, either in the command line or in the properties file.

The `configcopy` utility generates its own log file, `configcopy.log`, which is 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 attempt 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.

In addition, to avoid errors and prevent unstable target configurations due to missing components or faulty cross-component references (for example, invalid ID references from connection configurations to data-store configurations), be sure to adhere closely to the instructions provided in the following procedure.

1. Ensure access to the Connection Management Web Service is enabled for both the source and target PingFederate servers (see [Configure application authentication](#)).

2. Determine which component configurations need to be copied, including plug-ins.

For example, connection configurations always reference either adapter or token-translator configurations (or both) and may reference data-store configurations. These are all separate configurations, and must be copied separately (unless they already exist at the target) in conjunction with copying connection configurations.

Server Settings, unless preconfigured at the target, also need to be copied over separately.

Provisioning settings may be copied separately as needed to update target connections.

3. Determine whether any configuration property files or other supporting files need to be copied.

4. Ensure necessary plug-in 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:

```
<pf_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 (see [Certificate management](#) on page 178).

Private keys are not copied from server to server (public certificates may be copied); however, you may use `configcopy` to upload keys/certificates to the target server.

Make note of 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, either do so—you can use a `configcopy` command for this—or use one of the following work-arounds to ensure that `configcopy` can contact both servers:

Either:

- (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`.


Or:

- Install the Issuer certificate for the PingFederate SSL certificate into the trust store for the Server JRE under which `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 Server JRE under which `configcopy` runs.

7. Create properties files for the necessary commands and associated command-parameter values needed to copy the required configurations and any additional files.

Refer to the REAME.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 may also use command-line parameters (see information earlier in this section).

8. If you are copying connections, override ID properties referencing adapter, data stores or other plug-in configurations, as needed (see [step 2](#)).

 **Important:** Ensure that the plug-in 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 discussion above 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 IdP sites. Administrators can use this tool 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 may get 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 internal data store PingFederate uses to maintain provisioning synchronization between the LDAP user store and the target service.

Note that the tool creates its own log file, `provmgr.log`, located in the directory `<pf_install>/pingfederate/log`. You can control settings for this log, as needed, in the file `provmgr.log4j2.xml`, located in the `<pf_install>/pingfederate/bin` directory.

The following table describes the available global and channel-specific command arguments.

Command argument	Description
Global options	
<code>--help</code>	Describes the available options. The help is also displayed if the command is run with no arguments.
<code>--show-channels</code>	Lists all channels in a table format, showing for each: <ul style="list-style-type: none"> • 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 (for example, 5000/12 means 5000 users and 12 dirty records) • Source (as LDAP URL) • Target code

Command argument	Description
<code>--show-nodes</code>	Shows all the provisioning-server nodes with their status and the last timestamp (applicable only when failover provisioning is configured in the <code><pf_install>/pingfederate/bin/run.properties</code> file).
<code>--force-node-backup</code>	Sets the provisioner mode to FAILOVER for the associated PingFederate server node.
Use with node number: <code>-n <node ID></code>	
Channel-specific options	
 Note: With each command, specify the channel with the <code>-c <channel-id-number></code> argument; for example:	
<pre>provmgr -c 1 --show-source</pre>	
	You can determine channel ID numbers by using the global command: <code>provmgr --show-channels</code>
<code>--reset-group-timestamp</code>	<p>Deletes the user-group timestamp, which forces the provisioner to process the provisioning group on the next cycle, even if the timestamp on that group did not actually change.</p> <p>Depending on your LDAP server and administrative practices, you may 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.</p> <p> Important: This option should seldom be needed if users are deactivated rather than deleted. If it is needed, you may wish to schedule it when other network activity is low.</p>
<code>--reset-attribute-sync</code>	<p>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.</p> <p> Important: This option should be needed rarely and may consume considerable network resources, depending on the number of users. If it is needed, you may wish to schedule it when other network activity is low.</p>
<code>--reset-values-hash</code>	<p>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.)</p> <p>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. Note, however, that users whose recorded timestamp is unchanged are <i>not</i> updated.</p>
<code>--reset-all</code>	<p>Equivalent to using all three of the arguments above.</p> <p> Important: This option should be needed rarely if ever and may consume considerable network resources, depending on the number of users. If it is needed, you may wish to schedule it when other network activity is low.</p>
<code>--show-dirty-records</code>	Lists all users or groups that have not been provisioned or updated at the service-provider site.
<code>--show-dirty-group-records</code>	List groups that have not been provisioned or updated at the service-provider site.
<code>--show-dirty-user-records</code>	List all users that have not been provisioned or updated at the service-provider site.
<code>--show-group</code>	Shows all internal database fields related to the specified user or group, including transitory mapping fields (fields waiting to be pushed to the service provider); for a user, shows all LDAP attributes retrieved from the directory server.
<code>--show-user</code>	
Use with:	

Command argument	Description
-u <provider name> Or: -g <LDAP GUID>	<p> Note: You can obtain user or group names and GUIDs for dirty records, as needed, using any of the <code>--show-dirty-*</code> options (described above).</p> <p>The LDAP GUID, if used and if it is binary, should be entered in hexadecimal format (as shown in log files).</p> <p>Examples:</p> <pre>provmgr.sh --show-user -u john@example.com provmgr.sh --show-user -g ffd448643f812b43a0bee2504173f0</pre>
--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.
--delete-dirty-user-records	Removes all dirty user records from the internal store.
--delete-all	<p>The delete-all parameter removes all users and groups from the internal store and deletes the provisioning group timestamp and the last attribute-sync timestamp. The delete-all-users parameter deletes users and timestamps but retains groups.</p> <p>The effect of either command is to reset the channel to its initial state for user provisioning. All user metadata is lost and provisioning for the channel will start from the beginning, picking up all users (and groups if deleted) and pushing them to the service provider when the synchronization frequency interval (defined on the Server Configuration > Server Settings > Outbound Provisioning screen) has expired.</p> <p> Important: These options should be needed rarely if ever. If needed, you may wish to schedule the operation when other network activity is low.</p>
--delete-all-users	
--show-target	Displays the target configuration.
--show-source	Displays all source LDAP configuration parameters, including settings and location.


Customizable user-facing screens

PingFederate® supplies HTML templates to provide information to the end users or to request user input when processing their requests. These template pages are located in the `<pf_install>/pingfederate/server/default/conf/template` directory. They utilize the Velocity template engine, an open-source Apache project. For 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. Cascading style sheets 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
<code>\$escape</code>	A utility class that can be used to escape String variables inserted into the template, for example, <code>\$escape.escape(\$client.name)</code> where <code>\$client.name</code> is a second variable available for the page.
<code>\$HttpServletRequest</code>	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 (for example, the <code>User-Agent</code> HTTP header).
<code>\$HttpServletResponse</code>	A Java object instance of <code>javax.servlet.http.HttpServletResponse</code> . Used to modify the response in the template (for example, setting additional browser cookies).
<code>\$locale</code>	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.
<code>\$PingFedBaseURL</code>	The PingFederate base URL, as defined on the Server Configuration > Server Settings > Federation Info screen.
<code>\$templateMessages</code>	Used to localize messages in the template (based on user's Locale), an instance of <code>com.pingidentity.sdk.locale.LanguagePackMessages</code> . For more information, see the Javadoc for the <code>LanguagePackMessages</code> class in the directory <code><pf_install>/pingfederate/sdk/doc</code> .
<code>\$TrackingId</code>	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 an SSO error occurs during 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 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 **IdP Configurations** or **SP Configurations** menu.

The Velocity templates retrieve titles and other text from a message-property file, `pingfederate-messages.properties`, located in the `<pf_install>/pingfederate/server/default/conf/language-packs` directory. You may also localized these messages using the PingFederate localization framework.

 **Note:** If you have a clustered PingFederate environment, copy the customized (and localized) templates to each node.

IdP user-facing pages

HTML Form Adapter

Page title and <i>template file name</i>	Purpose	Type	Action
Sign On <i>html.form.login.template.html</i>	Displays a configurable user sign-on form when an HTML Form Adapter instance is in use for a connection.	Normal	User input required
Change Password <i>html.form.change.password.template.html</i>	Displays when a user attempts to change his or her password via the HTML Form Adapter.	Normal	User input required

Page title and <i>template file name</i>	Purpose	Type	Action
Change Password <i>html.form.message.template.html</i>	Displayed when a user has successfully changed his or her password.	Normal	User input required
Password Expiring <i>html.form.password.expiring.notification.template.html</i>	Displays a warning to an authenticated user if the password associated with the account is about to expire soon.	Normal	User input required
Password Management System Message <i>html.form.message.template.html</i>	Displayed when a user is being redirected to a password management system to change his or her password.	Normal	User input required
Forgot Password <i>forgot-password.html</i>	Displayed when a user attempts to reset his or her password via the HTML Form Adapter. If the user has entered a username in the sign-on form, the username is carried over to this form; otherwise, the user must enter his or her username to begin the self-service password reset process.	Normal	User input required
Forgot Password <i>forgot-password-resume.html</i>	Displays to prompt a user to enter the one-time password sent through an email or text message notification or to notify a user to refer to the email notification for password reset instructions. This template is applicable when the password reset type is Email One-Time Link , Email One-Time Password , or Text Message for the invoking HTML Form Adapter instance.	Normal	User input required
Forgot Password <i>forgot-password-change.html</i>	Displays to prompt a user to define a new password.	Normal	User input required
Forgot Password <i>forgot-password-success.html</i>	Displayed when a user has successfully reset his or her password.	Normal	User input required
Forgot Password <i>forgot-password-error.html</i>	Displayed when a password reset attempt fails.	Error	None
Unlock Your Account <i>account-unlock.html</i>	Displayed when a user has successfully unlocked his or her account through the HTML Form Adapter. This page also prompts the user to retain the current password or reset it.	Normal	User input required

Page title and <i>template file name</i>	Purpose	Type	Action
Security Question <i>html.form.login.challenge.template.html</i>	Displays a configurable challenge form for two-step authentication. This template can be used, for example, to create a RADIUS challenge form when using the RADIUS Username/Password Credential Validator.	Normal	User input required
User Consent <i>consent-form-template.html</i>	Displayed when a request requires a user's consent for an SSO to an SP.	Normal	User input required
Logout Confirmation <i>idp.slo.confirm.page.template.html</i>	Displayed when a user initiates a logout request. Applicable only if such confirmation is required, as configured on the IdP Configuration > Default URL screen.		
Sign Off <i>idp.logout.success.page.template.html</i>	Displayed when a user has successfully signed off in a configuration where the Logout Path field is configured but the Logout Redirect field is not.	Normal	None

Kerberos Adapter

Page title and <i>template file name</i>	Purpose	Type	Action
Error <i>kerberos.error.template.html</i>	Displays an error page to provide standardized information to the end user when the authentication attempt fails.	Error	Consult log
(No title) <i>meta.refresh.template.html</i>	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 <i>template file name</i>	Purpose	Type	Action
Select Authentication System <i>sourceid-choose-idp-adapter-form-template.html</i>	Displayed when multiple authentication sources are applicable and no preference is submitted as part of the request.	Normal	User input required
Sign On Error <i>idp.sso.error.page.template.html</i>	Displayed when IdP-initiated SSO fails and no other SSO error landing page is specified.	Error	Consult log and web developer

Page title and <i>template file name</i>	Purpose	Type	Action
Sign Off Successful <i>idp.slo.success.page.template.html</i>	Displayed when an SLO request succeeds and no other SLO success landing page is specified.	Normal	None
Sign Off Error <i>idp.slo.error.page.template.html</i>	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 <i>template file name</i>	Purpose	Type	Action
Working . . . <i>sourceid-wsfed-http-post-template.html</i>	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. Note that this page is normally not displayed if Javascript executes properly.	Normal	None
Signing off. . . <i>sourceid-wsfed-idp-signout-cleanup-invisible-template.html</i>	WS-Federation and OpenID Connect client IdP sign-out processing page. Note that no HTML is rendered in the browser.	Normal	None
Sign Off Successful <i>sourceid-wsfed-idp-signout-cleanup-template.html</i>	Indicates user signed out of the IdP under the WS-Federation protocol and lists each successful SP logout, when applicable. Also displays when an OpenID Connect client sends a logout request to the <code>/idp/startSLO.ping</code> endpoint to initiate an Asynchronous Front-Channel Logout process.	Normal	None

SP user-facing pages

Account linking

Page title and <i>template file name</i>	Purpose	Type	Action
Link Your Account <i>LocalIdPasswordLookup.form.template.html</i>	Used to authenticate a user at the SP when an account link needs to be established.	Normal	None
Account Unlinked <i>TerminateAccountLinks.page.emplate.html</i>	Communicates a user's successful <i>defederation</i> operation.	Normal	None

Single sign-on and logout

Page title and <i>template file name</i>	Purpose	Type	Action
Select Identity Provider <i>sourceid-saml2-idp-selection-template.html</i>	The user requested SP-initiated SSO, but the 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 <i>sp.sso.success.page.template.html</i>	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 SP Configuration > Default URLs screen.	Error	Consult web developer or specify default URL
Sign On Error <i>sp.sso.error.page.template.html</i>	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 <i>sp.slo.success.page.template.html</i>	Displayed when an SLO request succeeds and no other SLO success landing page is specified.	Normal	None
Sign Off Error <i>sp.slo.error.page.template.html</i>	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 <i>template file name</i>	Purpose	Type	Action
Signed Off <i>sourceid-wsfed-sp-signout-cleanup-template.html</i>	Displays the user's sign-out status.	Normal	None
Unable to Authenticate <i>sourceid-wsfed-idp-exception-template.html</i>	Displayed when an authentication challenge fails during WS-Federation processing.	Error	Consult log and web developer

Either IdP or SP user-facing pages

Page title and <i>template file name</i>	Purpose	Type	Action
Sign On <i>AbstractPasswordIdpAuthnAdapter.form.template.html</i>	Challenges user for credentials when authentication can take place via HTTP Basic Authentication or an HTML form, depending on the operational mode.	Normal	User must sign on
Submit Form	Whenever the server posts a form, this template is used to auto-	Normal	None

Page title and <i>template file name</i>	Purpose	Type	Action
<i>form.autopost.template.html</i>	submit the form. If Javascript is disabled, the user is prompted to click a button to post the form manually. Note that this page is normally not displayed if Javascript executes properly.		
Multiple Sign-On Delay <i>speed.bump.template.html</i>	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 <i>general.error.page.template.html</i>	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 <i>generic.error.msg.page.template.html</i>	General error, with error code.	Error	Consult log and check configuration Contact Ping Identity Support if unresolved
Error <i>http.error.page.template.html</i>	Indicates that an HTTP error has occurred and provides the HTTP status code.	Error	Consult log Contact Ping Identity Support if unresolved
Page Expired <i>state.not.found.error.page.template.html</i>	Displayed when simultaneous SSO requests (from multiple tabs using the same PF cookie) cause a user session to be overwritten or deleted and remaining requests attempt to retrieve the state fail.	Error	None

OAuth user-facing pages

The PingFederate® OAuth AS provides two screens presented to end users (resource owners) during OAuth transactions, one requesting approval of the scope of protected resources requested and one providing a means of revoking persistent access grants. These screens can be customized and branded as needed.


Page title and <i>template file name</i>	Purpose	Message type	Action
Client Access <i>oauth.access.grants.page.template.html</i>	Provides a means for the end users (resource owners) to revoke persistent access grants.	Normal	User input required

Page title and <i>template file name</i>	Purpose	Message type	Action
Request for Approval <i>oauth.approval.page.template.html</i>	<p>Advises resource owners that their information is being requested by the identified OAuth client and provides for approval/disapproval.</p> <p>This page appears for Implicit or Authorization Code grant types, either one time only or repeatedly depending on the Reuse Existing Persistent Access Grants for Grant Types setting in the OAuth Settings > Authorization Server Settings screen.</p> <p>In addition, the OAuth client configuration provides an option to bypass this approval page entirely, as needed for trusted clients. When applicable, select the Bypass Authorization Approval check box in client configuration screen.</p>	Normal	User input required

Customizable email notifications

PingFederate® supplies HTML templates to notify the administrators and the end users for various events. These template files are located in the `<pf_install>/pingfederate/server/default/conf/template/mail-notifications` directory.


You can modify these email notification template files in a text editor to suit the particular branding and informational needs of your PingFederate installation. Each template contains variables and standard HTML. At runtime, PingFederate supplies values for the variables and generates the email messages from the template files.

 **Note:** If you have a clustered PingFederate environment, copy the customized templates to each node.


Email notification is optional. Unless otherwise stated, all events and their respective options, including the recipient information, are configurable on the **Server Configuration > Server Settings > Runtime Notifications** screen. In addition, email configuration must be completed on the **Runtime Notifications > Email Server Settings** screen.

Account management events

Email notification for account management events is configurable on the **Server Configuration > Account Management** screen.

 **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.

 **Important:** In order for an administrator to receive emails about account management events, an email address must be provided on the **Account Management > User Information** screen when creating or modifying the administrator's profile.

Email subject and <i>template file name</i>	Event	Action
PingFederate Notification Settings Change Notification <i>message-template-notifications.html</i>	<p>An administrator has turned off the Notify Administrator of Account Changes option.</p> <p>An email notification is sent immediately to all administrators.</p> <p>The message includes the username of the administrator who made the change.</p>	Ensure this change is approved and legitimate.
PingFederate Email Change Notification <i>message-template-email.html</i>	<p>An administrator's email address has been updated by another administrator.</p> <p>An email notification is sent immediately to the previous email address and the updated email address.</p> <p>The message includes the username of the administrator who made the change.</p>	Ensure this change is approved and legitimate.
PingFederate Password Change Notification <i>message-template-password.html</i>	<p>An administrator's password has been changed.</p> <p>An email notification is sent immediately to the administrator whose password has been changed.</p> <p>The message includes the username of the administrator who made the change.</p>	Ensure this change is approved and legitimate.

Certificate events

If PingFederate® is configured to send email notification for certificate events, for any self-signed certificates that the administrators have enabled automatic certificate rotation on the **Server Configuration > Signing & Decryption Keys & Certificates** screen, PingFederate also sends email notification when a new pending certificate is created and when it is activated. The notification is sent to the email address that has been configured for the certificate events on the **Runtime Notifications** screen.

Email subject and <i>template file name</i>	Event	Action
A PingFederate Certificate Is About to Expire <i>message-template-cert-warning.html</i>	<p>A certificate is about to expire.</p> <p>The email notification is sent based on settings defined on the Runtime Notifications screen.</p> <p>The message includes the details of the certificate and the connections associated with it.</p>	<p>Create a new certificate and work with the applicable partners to update the expiring certificate.</p> <p>If self-signed certificate is used for signing or decryption, consider enabling certificate rotation while creating the new certificate.</p> <p>For a clustered PingFederate environment, replicate the</p>

Email subject and <i>template file name</i>	Event	Action
A PingFederate Certificate Has Expired <i>message-template-cert-expire.html</i>	<p>A certificate expired.</p> <p>The email notification is sent upon the expiration of a certificate.</p> <p>The message includes the details and the connections associated with the certificate.</p>	<p>configuration using the administrative console.</p> <p>Create a new certificate and work with the applicable partners to update the expiring certificate.</p> <p>If self-signed certificate is used for signing or decryption, consider enabling certificate rotation while creating the new certificate.</p> <p>For a clustered PingFederate environment, replicate the configuration using the administrative console.</p>
A New PingFederate Certificate Has Been Created <i>message-template-cert-rotation.html</i>	<p>A new pending certificate has been created for signing or decryption.</p> <p>The email notification is sent when a new pending certificate is created.</p> <p>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.</p>	<p>Work with the applicable partners to update the expiring certificate.</p> <p>PingFederate supports providing metadata for Browser SSO connections.</p> <p>For a clustered PingFederate environment, replicate the configuration using the administrative console.</p>
A PingFederate Certificate Has Been Updated <i>message-template-cert-deactivation.html</i>	<p>A new certificate for signing or decryption is activated.</p> <p>The email notification is sent when the new certificate is activated.</p> <p>The message includes the details of the new certificate and the affected connections.</p>	<p>None unless the applicable partners have not been notified or configuration has not been replicated in a clustered PingFederate environment.</p>

HTML Form Adapter events

The HTML Form Adapter supports notifying the end users when their password has been updated through the adapter. Administrators may enable this optional feature by selecting the **Change Password Email Notification** check box on a per-instance configuration basis. The notifications are sent to the end users.

In addition, administrators can optionally enable email-based self-service password reset and customize the email notifications.

Email subject and <i>template file name</i>	Event	Action
Password Change Notification <i>message-template-end-user-password-change.html</i>	<p>An end-user's password has been changed successfully through an instance of the HTML Form Adapter.</p>	<p>The end users should contact their IT administrators if they</p>

Email subject and <i>template file name</i>	Event	Action
	The email notification is sent immediately.	have not made any attempts to change their password.
Password Reset <i>message-template-forgot-password-link.html</i>	<p>An end user has initiated a self-service password reset request.</p> <p>Applicable when the password reset type for the invoking HTML Form Adapter instance is set to Email One-Time Link.</p> <p>The email notification is sent immediately.</p>	The end users should contact their IT administrators if they have not made any attempts to reset their password.
Password Reset <i>message-template-forgot-password-code.html</i>	<p>An end user has initiated a self-service password reset request.</p> <p>Applicable when the password reset type for the invoking HTML Form Adapter instance is set to Email One-Time Password or Text Message.</p> <p>The email notification is sent immediately.</p>	The end users should contact their IT administrators if they have not made any attempts to reset their password.
Password Reset Completed <i>message-template-forgot-password-complete.html</i>	<p>An end-user's password has been reset successfully through an instance of the HTML Form Adapter.</p> <p>Applicable when the password reset type for the invoking HTML Form Adapter instance is set to any method (but None).</p> <p>The email notification is sent immediately.</p>	The end users should contact their IT administrators if they have not made any attempts to reset their password.
Account Unlocked <i>message-template-account-unlock-complete.html</i>	<p>A user account has been unlocked successfully through an instance of the HTML Form Adapter.</p> <p>Applicable when the Account Unlock option is enabled for the invoking HTML Form Adapter instance.</p>	The end users should contact their IT administrators if they have not made any attempts to unlock their account.
Password Reset Failed <i>message-template-forgot-password-failed.html</i>	<p>An attempt to reset the end-user's password has failed.</p> <p>The email notification is sent immediately.</p>	The end users should contact their IT administrators if they have not made any attempts to reset their password, or to seek help in resetting their password.

As needed, administrators can localized these email messages using the PingFederate localization framework.


SAML metadata update events

PingFederate® supports automatic reloading of SAML metadata, which streamlines the maintenance of SAML connections. When notification for SAML metadata update events is enabled on the **Runtime Notifications** screen, PingFederate sends the notifications to the email address that has been configured for the event on the same screen.

Email subject and <i>template file name</i>	Event	Action
Your \${CONNECTION_NAME} \${SP_IDP} Connection in PingFederate Has Been Updated <i>message-template-metadata-updated.html</i>	PingFederate has updated a connection based on the partner's SAML metadata.	For a clustered PingFederate environment, replicate the configuration using the administrative console.
Please Review Your Updated \${CONNECTION_NAME} \${SP_IDP} Connection in PingFederate <i>message-template-metadata-updated-out-of-sync.html</i>	PingFederate has updated 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. For a clustered PingFederate environment, replicate the configuration using the administrative console.
Your \${CONNECTION_NAME} \${SP_IDP} Connection in PingFederate Is Out of Sync <i>message-template-metadata-out-of-sync.html</i>	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. For a clustered PingFederate environment, replicate the configuration using the administrative console.
Your Metadata Couldn't Be Downloaded from \${METADATA_URL_NAME} <i>message-template-metadata-url-notification.html</i>	PingFederate failed to download the SAML metadata from the partner or could not validate the digital signature of the metadata.	Consult log. Verify and update the metadata URL and its corresponding verification certificate.
Your Metadata for PingFederate \${SP_IDP} Connection \${CONNECTION_NAME} Wasn't Found <i>message-template-metadata-url-entity-id-missing.html</i>	The partner's metadata URL did not return the expected metadata for a given connection.	Consult log. Verify and update the metadata URL.

Licensing events

When notification for licensing events is enabled, PingFederate® sends license expiry information to the recipient configured for the event on the **Runtime Notifications** screen.

 **Tip:** To check the details of your license, sign on to the administrative console, point to the user icon on the upper-right corner, then click **About**. The license summary is displayed in a pop-up browser window.

Note that if an expiration date is specified, the license expires at the beginning of the day.

Email subject and <i>template file name</i>	Event	Action
Your PingFederate License Is About to Expire <i>message-template-warn.html</i>	The current license is about to expire. The email notification is sent 60 days ahead of the expiration.	Contact sales@pingidentity.com to renew the license before the expiration.

Email subject and <i>template file name</i>	Event	Action
	Applicable to licenses without connection groups.	
Your PingFederate License Is About to Expire <i>message-template-group-warn.html</i>	One of the connection groups in the current license is about to expire. The email notification is sent 60 days ahead of the expiration.	Contact sales@pingidentity.com to renew the license before the expiration.
	Applicable to licenses <i>with</i> connection groups.	
PingFederate License Expiration Notification <i>message-template-grace.html</i>	The current license expired. The email notification is sent upon the expiration.	Contact sales@pingidentity.com to renew the license.
	Applicable to licenses without connection groups.	
Your PingFederate License Has Expired <i>message-template-group-grace.html</i>	One of the connection groups in the current license expired. The email notification is sent upon the expiration.	Contact sales@pingidentity.com to renew the license.
	Applicable to licenses <i>with</i> connection groups.	
PingFederate Has Stopped Processing Requests <i>message-template-shutdown.html</i>	The current license and the grace period (if any) had lapsed.	Contact sales@pingidentity.com to renew the license.
	Applicable to licenses without connection groups.	


Email server verification

You must complete the email configuration on the **Runtime Notifications > Email Server Settings** screen in order to enable email notification for any events. Although optional, it is recommend that you enter a test recipient and run a verification test on the configuration screen.

Email subject and <i>template file name</i>	Event	Action
Test Message from Your PingFederate Server <i>message-template-test.html</i>	To verify the email server settings are functional.	Ensure the verification email is received by the intended recipient. Consult log, as needed.

Customizable text message

If you have configured self-service password reset (and optionally self-service account unlock) to use the **Text Message** option, you may customize (and localized) 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.

Localize messages for end users

PingFederate® support localization for three types of end-user messages: on-screen messages, email messages, and text messages (SMS). The English contents for each message type are located in the `<pf_install>/pingfederate/server/default/conf/language-packs` directory.

Message type	Default message file
On-screen messages	<code>pingfederate-messages.properties</code>
Email messages	<code>pingfederate-email-messages.properties</code>
Text messages (SMS)	<code>pingfederate-sms-messages.properties</code>

Follow these steps to localize messages for your end users.


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 browser settings.

For example, to localize the user-facing screen 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 (for example from `pingfederate-messages.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 instance, to localize the text messages in Canadian French, renamed the translated copy to `pingfederate-sms-messages_fr_CA.properties`.

 **Note:** The capitalization and underscore usage may 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.

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. Refer to the template files for examples.

 **Tip:** 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.

Be sure to return the value to `true` when testing is complete. Note that restarting the server is required for changes to configuration files.

Locale overrides by cookies

For convenience, an administrator or web developer might want to provide end-users a means of overriding browser language preferences temporarily by setting cookies; for example, by 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 via a cookie named `pf-accept-language`. The cookie value must conform to guidelines defined under [IETF BCP 47](https://tools.ietf.org/html/bcp47) (tools.ietf.org/html/bcp47). For more information about the Java core method that PingFederate uses to parse the cookie, see [Locale.forLanguageTag\(String languageTag\)](https://docs.oracle.com/javase/7/docs/api/java/util/Locale.html#forLanguageTag(java.lang.String)) (docs.oracle.com/javase/7/docs/api/java/util/Locale.html#forLanguageTag(java.lang.String)).

Note that this locale-override behavior is the default implementation of the Java interface `LocaleOverrideService` defined in the PingFederate SDK (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 `language-packs` 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 via the `LanguagePackMessages` class available in the PingFederate SDK. An instance of this class is passed into every template file and available for use there. For information, refer to the Javadoc for the class in the `<pf_install>/pingfederate/sdk/doc` directory.

Configure password policy

PingFederate® applies a configurable policy to passwords, pass phrases, and shared secrets defined by the administrators in the administrative console. 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
 - JMX
 - Connection Management
 - SSO Directory Service
- Passwords used by instances of the Simple Username Password Credential Validator
- Passwords used for encrypting certificates exported with their private keys
- Pass phrases used by IdP Discovery
- Passwords used by administrative console credentials when native authentication is used



Note: Passwords external to PingFederate—passwords used by instances of the Data Stores, for example—are not subject to this password policy.

1. Edit the `<pf_install>/pingfederate/server/default/data/config-store/password-rules.xml` file.
2. Save the changes.
3. Restart PingFederate.

For a clustered PingFederate environment, perform these steps on the console node. No changes or restart of PingFederate is required on the engine nodes.

Manage expired OAuth persistent grants

Persistent grants can result from the following OAuth use cases:

- OAuth clients using the **Refresh Token** grant type in conjunction with either the **Authorization Code** or **Resource Owner Credentials** grant type.


If the use cases involve mapping attributes from authentication sources (IdP adapter instances or IdP connections) or Password Credential Validator (PCV) instances to the access tokens (directly or through persistent grant extended attributes), such attributes and their values are stored along with the persistent grants so that they can be reused when clients subsequently present refresh tokens for new access tokens.

- OAuth clients using the **Implicit** grant type *and* the **Reuse Existing Persistent Access Grants for Grant Types** check box is selected in the **OAuth Settings > Authorization Server Settings** screen.

If the use cases involve mapping attributes from authentication sources or PCV instances to the access tokens (directly or through persistent grant extended attributes), attribute values are obtained at runtime for each token request. No attributes or their values are stored with the persistent grants.

Persistent grants (and the associated attributes and their values, if any) remain valid until the grants expired or are explicitly revoked.

PingFederate® automatically removes expired persistent grants and the associated attributes once a day. The cleanup process removes 500 expired grants at a time until all expired grants are removed. If expired grants are growing rapidly, you can optionally increase the frequency of the cleanup process.

 **Note:** Increasing the frequency of the cleanup process, the number of expired grants to be removed per batch, or both adds more workload to your storage server. We recommend making changes gradually to observe the impact, if any.

1. Optional: Adjust the frequency of the cleanup process.
 - a) Edit the <pf_install>/pingfederate/server/default/data/config-store/timer-intervals.xml file.
 - b) Update the AccessGrantCleanerInterval value.
 - c) Save your change.
2. Optional: Adjust the number of expired grants to be removed per batch.
 - a) Edit a configuration file in the <pf_install>/pingfederate/server/default/data/config-store directory.

Storage platform	Configuration file
Database server	org.sourceid.oauth20.token.AccessGrantManagerJ
PingDirectory™	org.sourceid.oauth20.token.AccessGrantManagerL
Microsoft Active Directory	org.sourceid.oauth20.token.AccessGrantManagerL
Oracle Directory Server Enterprise Edition	org.sourceid.oauth20.token.AccessGrantManagerL

- b) Update the ExpiredGrantBatchSize value if the storage platform is a database server.

If you use an LDAP server for grant storage, add the ExpiredGrantBatchSize property to the associated configuration file; for example:

```
<?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>
</c:config>
```

(In this example, the new value is 400.)

- c) Save your change.
3. If you have made any changes, restart PingFederate.

For a clustered PingFederate environment, perform these steps on the console node because the cleanup process only runs on the console node. No changes or restart of PingFederate is required on the engine nodes.

Specify the domain of the PF cookie

PingFederate® identifies sessions by their respective PF cookie. By default, the PF cookie is set without domain information in the HTTP header; for example:

```
Set-Cookie: PF=zOv4xxmzDI2rx1TFBFy78X;Path=/;Secure;HttpOnly
```

As needed, you may configure PingFederate to return the Set-Cookie HTTP header with domain information.

1. Edit the `<pf_install>/pingfederate/server/default/data/config-store/session-cookie-config.xml` file.
2. Add a `cookie-domain` element; for 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 the **Server Configuration > Cluster Management** screen.

 **Note:** It is not necessary to restart PingFederate on any running engine node.

Once this change is activated, PingFederate includes domain information in its Set-Cookie HTTP header; for example:

```
Set-Cookie:
PF=aDfPx6uwbbWGFhwE6zEhEG; Path=/; Domain=.example.com; Secure; HttpOnly
```

Extend the lifetime of the PF cookie

PingFederate identifies PingFederate-sessions by their respective PF cookies. Some adapters, such as the HTML Form Adapter, also utilize the PF cookie to manage their adapter-sessions.

By default, the PF cookie is a session cookie. You may extend the lifetime of the PF cookie by making it a persistent cookie. Unlike session cookies, persistent cookies are saved to disk, enabling your web browser to reuse them when restarted.

1. Edit the `<pf_install>/pingfederate/server/default/data/config-store/session-cookie-config.xml` file.
2. Modify the `cookie-max-age` value. The default value (-1) makes the 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 the **Server Configuration > Cluster Management** screen.

 **Note:** It is not necessary to restart PingFederate on any running engine node.

Configure forward proxy server settings

You can configure PingFederate[®] to send web traffic (HTTP, HTTPS, or both) that it initiates through a forward proxy server.

1. Edit the `<pf_install>/pingfederate/bin/run.properties` file.
2. Look 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; for example:

```
*.example.com|localhost
```
- 6. Save your changes.
- 7. Restart PingFederate.


For a clustered PingFederate environment, repeat these steps on each node.

Add 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.

1. Edit the `response-header-admin-config.xml` file or the `response-header-runtime-config.xml` file (or both) in the `<pf_install>/pingfederate/server/default/data/config-store` directory.
2. Save the changes.
3. Restart PingFederate.

For a clustered PingFederate environment, perform these steps on the console node, and then click **Replicate Configuration** on the **Server Configuration > Cluster Management** screen.

 **Note:** For each engine node, restart PingFederate to reload the `response-header-admin-config.xml` or `response-header-runtime-config.xml` file (or both) after the configuration is replicated.

Configure validation for the AudienceRestriction element

For any IdP connections 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.

As needed, you may disregard this validation condition on a per-connection basis.

1. Edit the `org.sourceid.saml20.util.VirtualIdentityUtil.xml` file in the `<pf_install>/pingfederate/server/default/data/config-store` directory.
2. For each IdP connection that you want to disregard the validation condition, add its **Partner's Entity ID** value as an entry inside the `c:map` element; for example:

```
<?xml version="1.0" encoding="UTF-8"?>
<c:config xmlns:c="http://www.sourceid.org/2004/05/config">
  <c:map name="AllowAnyVirtualServerIdInAudience">
    <c:item name="idp.example.com"/>
    <c:item name="sso.example.info"/>
  </c:map>
</c:config>
```

In this example, the first entry adds the IdP connection with a **Partner's Entity ID** of `idp.example.com` to the list 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 `sso.example.info`.

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 the **Server Configuration > Cluster Management** screen.



Note: It is not necessary to restart PingFederate on any running engine node.

Customize the heartbeat message

The heartbeat endpoint (`/pf/heartbeat.ping`) returns an "OK" browser message and an HTTP 200 status indication if the PingFederate server is running. You can customize the message by modifying a PingFederate® property and a Velocity template file.



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.

1. Set the `pf.heartbeat.system.monitoring` property to `true` in the `<pf_install>/pingfederate/bin/run/properties` file.
2. Restart PingFederate.
3. Edit the `<pf_install>/pingfederate/server/default/conf/template/heartbeat.page.template` Velocity template file to specify the desired information to be returned by the heartbeat endpoint. (An inline sample is provided. Template customization does not require a restart of PingFederate.)

For a clustered PingFederate environment, repeat these steps on each engine node.

Customize the favicon for application and protocol endpoints

PingFederate® provides a favorite icon (favicon) for its application (such as `/idp/startSSO.ping`) and protocol endpoints (for example, `/idp/SSO.saml2`).

1. Replace the `favicon.ico` file in the `<pf_install>/pingfederate/server/default/conf/template/assets/images` directory.
2. Restart PingFederate.

For a clustered PingFederate environment, repeat these steps on each engine node.

Configure the behavior of searching multiple data stores with one mapping

Starting with PingFederate® 8.1, if a data store uses results from previous queries as input, and if the previous queries return no result, PingFederate continues the SSO process by moving on to the next data store in the setup. If you prefer PingFederate to abort the SSO requests (the default behavior of PingFederate 8.0 and older), you can override the behavior by modifying a configuration file.


1. Edit the `org.sourceid.saml20.domain.AttributeMapping.xml` file in the `<pf_install>/pingfederate/server/default/data/config-store` directory.

Create this file if it does not exist.

2. Change the value of the `AbortOnAttrLookupFailure` element from `false` (the default value) to `true`.

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>
</c:config>
```

 **Note:** Removing the `org.sourceid.saml20.domain.AttributeMapping.xml` file from `<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 the **Server Configuration > Cluster Management** screen.

 **Note:** It is not necessary to restart PingFederate on any running engine node.


For more information, see [Attribute mapping with multiple data sources](#) on page 506.

System settings

The System Settings links on the Server Configuration menu provide access to global settings that may apply to either an IdP or an SP configuration.

This chapter covers:

- [Server settings](#) on page 145
- [Manage data stores](#) on page 158
- [IdP discovery](#) on page 174
- [Configure redirect validation](#) on page 176

 **Note:** The information in this chapter is presented from the viewpoint of an administrative user with “Admin” permissions (see [Account management](#) on page 108).

Server settings

Server settings include unique federation server identifiers, the designation of your site's federation role (SP, IdP, or both), and your enabled federation protocols (see [Supported Standards](#) chapter of the *Get started with PingFederate* guide).

Server settings also include system-administration configuration (one-user or multi-user), email notification options and setup, and a shortcut link to account management (when multi-user administration is enabled).

If you have enabled Auto-Connect, Outbound Provisioning, or both, you must configure several parameters specific to those features in the System Settings task flow.

You configure many of these settings initially during the installation setup, but you can change or add to them as needed from the Main Menu.

 **Note:** For information about WS-Trust STS Settings, see [Configure STS authentication](#) on page 400.

Information in this section covers:

- [Set administration options](#) on page 146
- [Enter system information](#) on page 146
- [Configure runtime notifications](#) on page 146
- [Configure runtime reporting](#) on page 147
- [Manage PingOne settings](#) on page 150
- [Managing accounts](#) on page 151
- [Choose roles and protocols](#) on page 151
- [Specify federation information](#) on page 152
- [Configure system options](#) on page 153
- [Configure outbound provisioning settings](#) on page 155
- [Configure metadata signing](#) on page 156
- [Configure metadata lifetime](#) on page 157
- [Edit and save server settings](#) on page 157

Set administration options

On the System Administration screen, PingFederate provides a choice of single- or multi-user access to the administrative console.



Note: If you are using your network's LDAP user-data store or client certificates for administrative-console authentication, this screen is not presented (see [Alternative console authentication](#) on page 112).

If you choose Single-user Administration, the console is accessible only by using the default Administrator ID, for which full privileges are provided. Multi-user Administration (the default) provides role-based access control (see [Account management](#) on page 108).



Tip: To return to single-user administration after having previously enabled multi-user, only one user can be marked as active under Account Management.

Enter system information

On the System Info screen, you provide general information about your company.

To reach this screen

1. Click **Server Configuration** on the Main Menu.
2. Click **Server Settings** under System Settings.
3. Click **System Info** on the Summary screen.

Configure runtime notifications

Depending on your licensing agreement, your PingFederate license may have an expiration date. On the **Server Configuration > Server Settings > Runtime Notifications** screen, when the **Notification for server licensing events** check box is selected, PingFederate sends an email warning to an email recipient (or an email group address) when your license is about to expire. Note that this option does not appear when you have a perpetual license.

You can also configure the server to send an email notification for certificate events to the same or a different recipient. When the **Notification for certificate events** is selected, PingFederate sends email notifications when a certificate used by PingFederate is about to expire or has expired. In addition, for self-signed certificate that you have enabled automatic certificate rotation, PingFederate also sends email notifications when a new key pair and a new certificate are created, and when the new certificate is activated.

Finally, you can also configure PingFederate to send email notifications when it detects changes for SAML connections that are configured to pull metadata periodically from the partners. Select the **Notification for SAML metadata update events** check box and enter an email address for the same or yet another recipient.



Tip: In a clustered PingFederate environment, the email notifications are only sent by one of the nodes in the cluster: the console node or any engine node. If this node leaves the cluster (planned, or not), another node picks up this task automatically.

Regardless of runtime notification settings, PingFederate logs license information, certificate events, and SAML metadata update events in the server log by default (see [PingFederate log files](#) on page 95).

To configure notifications:

1. Select the check box next to the type of notification you want, and enter an email address.
2. If you have selected the **Notification for certificate events** check box, enter a warning time period in the **Initial Warning Event** field (optional) and in the **Final Warning Event** field.




Note: When PingFederate is configured to send notifications for certificate events, it also sends notification if a license expires.

3. If you have not previously configured PingFederate to access your email server, click **Email Server Settings** (see [Manage email configuration](#) on page 114).

Configure 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).

 **Tip:** You can also use HTTP requests at any time to verify the status of the PingFederate server (see [Customize the heartbeat message](#) on page 144).

Additionally, you may 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 (see [Security audit logging](#) on page 100).

Configure SNMP monitoring

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.

Gets

PingFederate responds to two SSO and 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

In addition, because PingFederate is built within an existing Jetty framework, *Gets* include a variety of server information available via 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 148).

 **Note:** Some operating systems (Solaris 10, for example) may not allow the SNMP agent to bind to privileged ports: those below 1024. Consult your operating system's documentation on how to get 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.

Configure access to SNMP monitoring on the **Server Configuration > Server Settings > Runtime Reporting** screen. As needed, you may configure both *Gets* and *Traps*.

- Optional: 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.
- Optional: 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 and verify the result.

Runtime monitoring using JMX

Similar to SNMP, 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 (for example, the standard Server JRE client JConsole) may request runtime information from the PingFederate® JMX server.

! **Important:** Authentication is required for JMX-client access to PingFederate runtime data (see [Configure application authentication](#)).

PingFederate JMX server reports monitoring data for SSO and SLO transactions. In addition, as with SNMP monitoring, numerous Jetty-standard MBeans are available to the PingFederate server's JMX clients.

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 via JMX, that administrators may find useful to supplement information provided via the PingFederate-specific MBeans.


MBean	Attributes
org.eclipse.jetty.server.connectorstatistics (For Jetty connectors including the primary and secondary PingFederate runtime server ports)	<p>connections – The total number of TCP connections accepted by the server.</p> <p>connectionsDuration* – How long connections are kept open. Maximum, mean, standard deviation, and total accumulated time are available.</p> <p>connectionsOpen – The current number of open connections. Maximum is also available (connectionsOpenMax).</p>
org.eclipse.jetty.server.handler.statistics	<p>requestsHandled – Total number of requests received.</p> <p>requestsActive – Number of requests currently being processed. Max is also available.</p> <p>requestTime – Request duration. Maximum, mean, standard deviation, and total accumulated time are available.</p> <p>responses1xx, responses2xx, responses3xx, ... – Total number of requests that returned HTTP status codes of 1xx, 2xx, 3xx, etc.</p>
org.eclipse.jetty.util.thread.queuedthreadpool (Two pools: one for the runtime server, with 200 maximum threads; one for the administrative console, with 20 maximum threads)	<p>idleThreads – Number of idle threads currently available.</p> <p>threads – Number of threads currently running (including both idle and active).</p> <p>minThreads – Minimum number of threads in the pool.</p> <p>maxThreads – Maximum number of threads in the pool.</p> <p>lowOnThreads – A boolean flag indicating whether the pool is running low on threads.</p>
java.lang: Memory	Various attributes measuring CPU usage and memory.
java.lang: MemoryPool	
java.lang: GarbageCollection	

MBean	Attributes
-------	------------


java.lang: OperatingSystem	
-------------------------------	--

Advanced JMX configuration

PingFederate uses port 1099 for its JMX server. To change the port or other JMS configuration items, if needed, modify the `jmx-remote-config.xml` configuration file in the `<pf_install>/pingfederate/server/default/conf` directory.


 **Note:** When connecting to the JMX service using SSL (the default), ensure that the client trusts the PingFederate SSL server certificate presented (see [Manage SSL server certificates](#) on page 179).

Monitor outbound provisioning

 **Important:** Monitoring of outbound provisioning transactions using JMX has been deprecated. PingFederate® now logs outbound provisioning events to `provisioner-audit.log` (see [Outbound provisioning audit logging](#) on page 102).

Follow these steps to re-enable JMX monitoring.

1. Edit the `<pf_install>/pingfederate/bin/run.properties` file.
2. Change the value of the `provisioner.events.monitor` property from `LOG` to `JMX`.
3. Edit the `<pf_install>/pingfederate/server/default/conf/jmx.remote.access` file.
4. Change the value of the `default-jmx-principal` property from `readonly` to `readwrite`.
5. Save all changes.
6. Restart PingFederate.

 **Note:** Repeat these steps on each PingFederate server configured to process outbound provisioning.

When JMX monitoring for outbound provisioning is enabled, PingFederate provides an MBean called `pf.provisioning:type=saas.provisioning.events`. The MBean exposes five JMX operations, each corresponding to the Java methods described in the following table. Each method returns a `CompositeData` object, which allows for the retrieval of complex data without requiring application-specific code to reside with the JMX client.

Method	Description	Parameter
<code>viewEvents(Boolean wasSuccessful, String eventTypeStr, String fromDate, String toDate)</code>	Gets an array of specific events based on the given criteria. The parameters filter the data collectively; that is, they are joined logically by “and”.	<p><code>wasSuccessful</code> – If <code>true</code>, returns information only on successful transactions; <code>false</code> returns information only on failed transactions; <code>null</code> returns all transactions.</p> <p><code>eventTypeStr</code> – The type of event. Valid values are: <code>CREATE</code>, <code>UPDATE</code>, <code>DISABLE</code>, <code>ENABLE</code>; <code>null</code> or an empty string returns all types.</p> <p><code>fromDate</code> – See the note at the end of this table.</p> <p><code>toDate</code> – See the note at the end of this table.</p>
<code>eventSummaryReport(String fromDate, String toDate)</code>	Gets a summary of transactions counts for the given time period. Counts are provided for success,	See the note at the end of this table.

Method	Description	Parameter
	failure, and total. Each count includes a drill-down capability, providing counts by event type.	
<code>provisioningCycleSummary (String fromDate, String toDate)</code>	Gets a total count of provisioning cycles for the given time period. The drill-down provides information for each cycle, including success totals for users and groups added, modified, or removed in the internal tracking database; for the target data store, success and failure totals are listed for each type of transaction.	See the note at the end of this table.
<code>eventSummaryReportAllData ()</code>	Gets a summary of transaction counts with no time constraints (equivalent to <code>eventSummaryReport</code> with <code>null</code> or empty strings used as parameters).	None.
<code>eventSummaryRollup ()</code>	Gets a report representing an aggregate of multiple Summary Reports covering the last 0, 1, 2, 7, 30, 60, 90, 180, and 360 days.	None.



Note: Date parameters may be formatted as either `yyyy`, `yyyy-MM-dd`, or `yyyy-MM-dd HH:mm:ss`. A `null` value or an empty string for a date parameter indicates no constraint for that end of the range.

Manage PingOne settings

After connecting PingFederate® to PingOne® through a managed SP connection, use the **Server Configuration > Server Settings > PingOne Settings** screen to manage the following options:

- [SSO from PingOne to the PingFederate administrative console](#) on page 150
- [Monitoring of PingFederate from PingOne](#) on page 151

PingFederate also automatically rotates the signing certificate used by the managed SP connection, see [Managed SP connection to PingOne and signing certificate](#) on page 188.

SSO from PingOne to the PingFederate administrative console

In PingFederate® 8.0 (or higher), if you have selected to connect PingFederate to PingOne® during the initial setup process, the option to SSO from PingOne to the PingFederate administrative console is enabled for you.



Tip: In this scenario, the initial setup wizard does not create any local administrative login. If you subsequently decide to disable the SSO from PingOne option, the administrative console will bring you to the **Account Management** screen and prompt you to create at least one user with the User Admin role for later use.


If you setup PingFederate without PingOne at the beginning but have subsequently created a managed SP connection to PingOne using the **Connect to PingOne** wizard, use the **Server Configuration > Server Settings > PingOne Settings** screen to enable this option.

Additionally, you can continue to sign on to the administrative console via native or alternative console authentication using the direct login page. You can also disable the direct login page to enforce the policy that administrators must SSO to the administrative console using their PingOne credentials.

- To SSO to the administrative console:
 - a) Verify the feature is enabled in the **PingOne Settings** screen.
 - b) Launch your browser and browse to:

`https://<DNS_NAME>:9999/pingfederate/app`

where `<DNS_NAME>` is the fully qualified name of the PingFederate server.

 **Tip:** To change the default administrative port number, see [Configure PingFederate properties](#) on page 93.

If you have already logged on to the PingOne admin portal, the PingFederate administrative console becomes available. If you are not logged on to the PingOne admin portal, you will be prompted by PingOne to enter your admin portal credentials. Upon verification, the PingFederate administrative console becomes available.


- To sign on via native or alternative console authentication:

- a) Launch your browser.
- b) Browse to:

`https://<DNS_NAME>:9999/pingfederate/app?service=page/directLogin`

- 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.

After restart, the direct login page is disabled. Administrators can only SSO to the PingFederate administrative console from PingOne at `https://<DNS_NAME>:9999/pingfederate/app`.

If you wish to re-enable native or alternative console authentication, update the `pf.console.authentication` property accordingly and restart PingFederate.

Monitoring of PingFederate from PingOne


After establishing a managed SP connection to PingOne®, you can monitor PingFederate® from the PingOne admin portal.

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 connected/disconnected status with the date last seen. For each server, you can also drill down for additional information such as CPU load, JVM memory information, and system memory information.

If you do not wish to monitor PingFederate from the PingOne admin portal, clear the option check box and click **Save** on the **Server Configuration > Server Settings > PingOne Settings** screen.

Managing accounts

When you choose multi-user system administration, you can create users during installation or while configuring Server Settings (see [Set administration options](#) on page 146).

 **Note:** If you are using your network's LDAP user-data store or client certificates for PingFederate authentication, the Account Management screen is not presented (see [Alternative console authentication](#) on page 112).

Alternatively, you can set up and maintain user accounts later as a separate task (assuming you have user administration permissions—see [Account management](#) on page 108). By default for installation, the user “Administrator” has full system permissions.


- To continue, click **Next** or **Save**.
- For information about adding or managing users, see [Account management](#) on page 108.

Choose roles and protocols


On the **Roles and Protocols** screen, select the roles your organization plays and the sets of standards you will use with your PingFederate server. Depending on the selected roles and protocols, you may be prompted to provide

additional information in a subsequent screen. If your use cases require roles or protocols that have not yet been selected, you must return to this screen to make the selections before you can configure those new use cases.

1. Go to the **Server Configuration > Server Settings > Roles & Protocols** screen.
2. Select your federation roles, and then select the applicable protocols.


 **Note:** Outbound provisioning for SaaS applications requires the use of the SAML 2.0. (For more information, refer to the *Quick Connection Guide* contained in the PingFederate SaaS Connector package for your service provider.)

3. Optional: If you are using SAML 2.0 and want to configure Auto-Connect, select that feature for your roles.


 **Note:** Clearing this check box does *not* deactivate an existing Auto-Connect configuration in production. If you have already deployed Auto-Connect and wish to suspend the deployment for any reason, use the **Initial Setup** summary screens for your respective role.

4. Optional: If you are using PingFederate as an IdP for provisioning or have installed a SaaS Connector package, select the **Outbound Provisioning** check box.

If such check box is not available, verify that your PingFederate license includes the **Outbound Provisioning** capability and the outbound provisioning properties are configured in the `<pf_install>/pingfederate/bin/run.properties` file.

 **Note:** After provisioning is configured for a connection, you cannot clear this check box—you must delete all provisioning configurations first. To suspend provisioning for an SP partner, you can deactivate the specific configuration. Alternatively, you can deactivate the associated SP connection; note, however, that this will also disable SSO/SLO transactions.

5. Optional: If you are using PingFederate as an SP for provisioning, select the **Inbound Provisioning** check box.
6. Optional: If you are using SAML 2.0 XASP as an SP for multiple IdP connections, you may select the option to determine dynamically which connection to use, based on the X.509 certificate presented.

 **Tip:** After you make this selection and create XASP IdP connections, configure dynamic IdP discovery via the **Attribute Requester Mapping** link on the **SP Configuration** menu. Once the mapping is configured, you cannot clear the check box on the **Roles and Protocols** screen unless you first delete the mapping.

For general information about XASP, see [Attribute Query and XASP](#) on page 43 in the “Supported Standards” chapter of the *Get started with PingFederate* guide.


7. Click **Next** (or **Save**, if you are modifying existing selections).

For information about configuring settings associated with your selections, see these relevant topics:

- [OAuth configuration](#) on page 241
- [Identity provider SSO configuration](#) on page 275
- [Outbound provisioning for IdPs](#) on page 81
- [Service provider SSO configuration](#) on page 330
- [Provisioning for SPs](#) on page 82
- [WS-Trust STS configuration](#) on page 400
- [IdP discovery](#) on page 174
- [Using Auto-Connect](#) on page 79

Specify federation information

This information identifies your federation deployment to your partners, according to the protocol(s) you support.

 **Note:** You must provide an ID that uniquely identifies your federation gateway for each protocol you support. For WS-Trust STS, IDs are required for both SAML 2.0 and SAML 1.x, regardless of browser-based 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 (see [Federation Server Identification](#)).

You can also use a different ID for Auto-Connect transactions (see [Using Auto-Connect](#) on page 79).

Field descriptions

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 (see Export metadata to an XML file on page 116).
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: <code>pingidentity.com</code> . The SAML SourceID used for artifact resolution is derived from this ID using SHA1.
Auto-Connect Entity ID	(Optional) If you are using Auto-Connect, you can specify a unique ID here for Auto-Connect processing. The value must be a fully qualified URL and should match the CN of your Auto-Connect certificates (see Auto-Connect security model on page 80). When a value is supplied, this ID is used instead of the SAML 2.0 Entity ID in your server's Auto-Connect metadata, as well as in associated SSO/SLO requests and responses. Use this field if you have configured regular, static SAML 2.0 connections to other partners and your SAML 2.0 Entity ID is not a fully qualified URL (see Using Auto-Connect on page 79).
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.



Note: The fields available on this screen depend on the federation protocols enabled on your server (see [Choose roles and protocols](#) on page 151).

To reach this screen:

1. Click **Server Configuration** on the Main Menu.
2. Click **Server Settings** under System Settings.
3. Click **Federation Info** on the Summary screen.

Configure system options


The **Server Configuration > Server Settings > System Options** screen provides global settings that allow you to:

- Turn off automatic multi-connection error checking
- Define a caching interval for data-store validation
- Configure incoming proxy settings

Configure automatic connection validation

Automatic multi-connection error checking occurs whenever you access certain supporting components, such as the **Adapters** screen, the **Password Credential Validators** screen, and the **Identity Store Provisioners** screen. The intent is to verify that all configured connections have not been adversely affected by subsequent changes in the supporting components.


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. Doing so also allows you to make configuration changes without fixing dependency errors (if any).

 **Note:** When automatic connection validation is turned off, multi-connection error checking is deferred until you access the connection lists on the administrative console. While you are free to make configuration changes without being prompted to fix all dependency errors immediately, keep in mind that the configuration changes could potentially introduce service disruption.

For example, if you remove an extended attribute (from an IdP adapter instance) that has been configured as the source of the SAML_SUBJECT attribute in an SP connection, users will not be able to complete SSO requests through this SP connection until you reconfigure the connection. When you access the connection on the **Connections** screen, the administrative console displays a visual cue, indicating the error condition. To resolve the error, open the connection by selecting its name and follow the on-screen instructions to modify the configuration that requires your attention.

- On the **Server Configuration > Server Settings > System Options** screen, select the **Disable automatic connection validation** check box if you do not want PingFederate® to validate connection settings automatically.

(This check box is not selected by default.)

 **Note:** This option does not affect the validation of connections as they are being configured or modified. Moreover, individual connections are always validated automatically when you access them on the **Connections** screen, regardless of the setting on this screen.

Specify data-store validation intervals

Automatic data-store validation tests occur by default for any adapter-to-adapter mappings or partner connections. This validation test occurs at various points within the administrative console.


The **Data Store Validation Interval (secs)** field defines the length of time for which a successful data-store validation is cached. (Only successful test results are cached.) The data-store connection is validated using the cached result without performing the test, speeding up the validation process.

- On the **Server Configuration > Server Settings > System Options** screen, override the default value, as needed.

The default interval is five minutes (300 seconds). A value of 0 turns off the caching and validation tests are executed with each access.

Configure proxy options

When PingFederate® is deployed behind a proxy server or load-balancer, the options described in the following four sections enable PingFederate to use information in HTTP headers added by the proxy server to construct correct responses. These options, configurable on the **Server Configuration > Server Settings > System Options** screen, apply globally to all incoming requests.

 **Note:** These settings override Jetty proxy options.

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.

It is common for proxies to 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 into 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. It is common for proxies to 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 are using mutual client certificate authentication for either WS-Trust STS authentication or SAML back-channel authentication and would like 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 in the **System Options** screen.

For example, suppose you have configured the Apache HTTP Server 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 CHAIN0 "%{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. In addition, the specified values must match the header names used in the Apache HTTP Server configuration (with the exception of 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 proxy server are made over HTTPS, even if HTTP is used between the proxy server and PingFederate.


Configure outbound provisioning settings

On the **Server Configuration > Server Settings > Outbound Provisioning** screen, select the database that PingFederate® should use internally to facilitate provisioning for service providers when PingFederate is configured as an IdP.

The database stores the state of synchronization between the source data store and the target data store, enabling periodic checking to determine whether updates are required at the target site. PingFederate checks the source data

store for changes every minute by default. As needed, you may change the provisioning synchronization frequency on this screen as well.

 **Caution:** A pre-installed, default HyperSQL database is selected for initial setup and testing. However, we strongly recommend that you choose your own, secured database for production deployments.

 **Note:** PingFederate has been tested using HyperSQL, Microsoft SQL Server, Oracle Database, Oracle MySQL, and PostgreSQL as internal provisioning data stores. However, any relational database should work as well; adaptable setup scripts used for HyperSQL, Microsoft SQL Server, Oracle Database, Oracle MySQL, and PostgreSQL are provided in the `<pf_install>/pingfederate/server/default/conf/provisioner/sql-scripts` directory.

Note that the **Outbound Provisioning** screen appears only when the **Outbound Provisioning** role is enabled on the **Server Configuration > Server Settings > Roles & Protocols** screen.

1. Select a data store from the list.

If the data store 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 data store.

2. Optional: Change the **Synchronization Frequency** value.

The default value is 60 (seconds).

Configure metadata signing

PingFederate generates publicly available metadata for partners through the federation metadata endpoint (`/pf/federation_metadata.ping`). Although optional, it is recommended to sign the metadata, such that partners can verify the authenticity of the metadata.

1. Go to the **Server Configuration > Server Settings > Metadata Signing** screen.
2. Select a certificate from the **Signing Certificate** list.


 **Important:** If you use Auto-Connect, the certificate CN must match the domain name associated with the **Auto-Connect Entity ID** field in the **Server Configuration > Server Settings > Federation Info** screen.

If you have not yet created or imported your certificate into PingFederate, click **Manage Certificates** and use the **Certificate Management** workflow 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.

 **Note:** For Auto-Connect, the metadata signing certificate must be trusted by your partner.

To specify a certificate:


1. Select the certificate from the drop-down list.

If you have not yet created or imported your certificate into PingFederate, click **Manage Certificates** (see [Manage digital signing certificates and decryption keys](#) on page 184).

2. Optional: Select the Signing Algorithm from the drop-down list.

The default selection is **RSA SHA256** or **ECDSA SHA256**, depending on the Key Algorithm value 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.

 **Note:** For Auto-Connect, the metadata signing certificate must be trusted by your partner (see [Auto-Connect security model](#) on page 80).

Configure metadata lifetime

PingFederate provides metadata for SAML and WS-Federation connections and supports auto-update for SAML connections by reloading metadata URLs provided by the partners.

Metadata publication

PingFederate includes expiration information in metadata. This metadata expiration ensures that partners always have reasonably up-to-date information about your server. You may elect to use the default time period (one day) or adjust the **Cache Duration** value on the **Server Configuration > Server Settings > Metadata Lifetime** screen.

Partners using Auto-Connect metadata also cache the metadata for the future requests during the lifetime of the metadata. After the lifetime has expired, the metadata is retrieved again.

Metadata consumption

PingFederate supports automatic reloading of metadata by URL for SAML connections. It checks daily by default. As needed, you can tune the frequency by modifying the **Reload Delay** value on the Metadata Lifetime screen. For more information about automatic reloading of SAML metadata by URL, see [Import SP metadata](#) on page 286 or [Import IdP metadata](#) on page 342.

Edit and save server settings


On the Server Settings Summary screen you can view, edit, and save your configuration.

- Click **Save** if you are finished with this configuration, or click any heading to make changes.

Connect to PingOne from Server Configuration

[PingOne](#) is Ping Identity's multi-tenant, identity-as-a-service (IDaaS) solution. PingOne® enables browser-based SSO and user provisioning for Identity Providers, and provides application providers with a rapid-deployment SSO capability. PingOne can be used together with PingFederate® to provide a powerful solution combining the benefits of an on-premise deployment with the flexibility of a cloud solution.



To leverage this unique capability, you need a PingOne account and a *managed* SP connection from PingFederate to PingOne. If you do not have a PingOne account, sign up at pingone.com.

 **Note:** A managed SP connection to PingOne is a connection created by the initial setup wizard or the **Server Configuration > Connect to PingOne** wizard in PingFederate 8.0 (or higher).

If an SP connection to PingOne was created in PingFederate 7.3 (or older), you can delete that connection and let the Connect to PingOne wizard create a new managed SP connection to PingOne for you. The benefits of a managed connection include the capability to SSO from PingOne to the PingFederate administrative console, monitoring of PingFederate from PingOne, automatic metadata exchange between PingFederate and PingOne, and automatic certificate rotations (see [Manage PingOne settings](#) on page 150 and [Managed SP connection to PingOne and signing certificate](#) on page 188).

1. Go to the **Server Configuration** menu.
2. Click **Connect to PingOne**.
3. On the **Identity Provider Configuration** screen:
 - Click **New Active Directory Configuration** if you need to establish a new LDAP data store to your Active Directory.
 - Click **Existing AD or Other User Store** if you have already created an LDAP data store to your Active Directory.
4. Complete the configuration.

Directory option	Process
New Active Directory Configuration	When selected, the Connect to PingOne wizard guides you through the process of: <ul style="list-style-type: none"> • Connecting your Active Directory as an LDAP data store

Directory option	Process
Existing AD or Other User Store	<ul style="list-style-type: none"> • Creating adapters and authentication selectors to authenticate end users via the Kerberos protocol or a login form based on your network topology and the browsers being used • Enabling users and group provisioning to PingOne • Enabling the capability to SSO from PingOne to the PingFederate administrative console <p>When selected, the Connection wizard initiates a new managed SP connection to PingOne for you. The attribute contract is automatically extended with the following six attributes:</p> <ul style="list-style-type: none"> • email • fname • lname • memberOf • objectGuid • userPrincipalName <p>Follow the step-by-step wizard to:</p> <ol style="list-style-type: none"> 1. Add one or more IdP adapter instances to the managed SP connection. 2. Select to retrieve additional attributes from your data store as needed. <ul style="list-style-type: none">  Tip: If you use the HTML Form Adapter created by the initial setup wizard (in PingFederate 8.0 or newer), you can map the six attributes from the adapter contract. No data store lookup is required.  Note: If you need to retrieve objectGUID from an LDAP data store, configure objectGUID to be handled as binary data in the LDAP data store and set Attribute Encoding Type to Hex on the LDAP Binary Attribute Encoding Types screen. <p>(For more information, see Specify LDAP binary attributes on page 164 and Specify encoding for LDAP binary attributes).</p> 3. Fulfill the attribute contract. 4. Configure outbound provisioning to PingOne. 5. Save the new managed SP connection to PingOne.

Manage data stores

PingFederate® can connect to data stores to retrieve user attributes for outbound SaaS or SCIM provisioning, Browser SSO, WS-Trust STS, and OAuth transactions. It can also read and write to data stores for inbound just-in-time or SCIM provisioning.

Browser SSO and WS-Trust STS

As an IdP, you may fulfill an attribute contract that requires information beyond that which can be derived from the user's session. For example, this information may include attributes such as an email address, a job title, or any data that can be used to customize a user's experience at the SP site.

As an SP, you may use data stores to retrieve additional attributes to package with the IdP's assertion data to meet SP adapter or token generator requirements. Such attributes may be needed, for example, to establish authorization levels or to manage the local account.

OAuth

As an OAuth authorization server (AS), you may pull user attributes from data stores to fulfill an access token attribute contract. If you support OpenID Connect, you may also map user attributes from data stores to an

OpenID Connect policy contract. For instance, you may want to return an email address and mobile phone number when an OAuth client connects to the UserInfo endpoint to retrieve additional attributes using an access token.

Outbound provisioning for IdPs

For outbound provisioning, you configure PingFederate to connect to two data stores:


- A data store from which PingFederate provisions users to the target SP.
- Another data store to monitor the state of the user store and to keep user data synchronized between the IdP and the target SP.

Inbound provisioning for SPs

For inbound provisioning, you connect to a data store to manage local user records based on inbound requests.


You can add data stores at any time. Standard data stores include JDBC-enabled databases and LDAP v3-compliant directories. Alternatively, you can develop a driver using the PingFederate Custom Source SDK to connect to non-JDBC databases.

1. Go to the **Server Configuration > Data Stores** screen.
2. In the **Manage Data Stores** screen, you can perform the following tasks:
 - To configure a new instance: click **Add New Data Store**, and then follow the wizard to configure a connection to a JDBC database, an LDAP directory server, or a custom data store.
 - 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**.

 **Important:** You must have current connectivity from PingFederate to a data store in order to create or modify the configuration. If you find that the configuration is not editable, then your connection has been lost due to a system problem not related to the PingFederate server. The problem must be identified and corrected before you can continue.

Configure a JDBC database connection

You configure access to a database by providing basic JDBC information.




 **Note:** PingFederate has been tested with vendor-specific JDBC 4.1 drivers. To obtain your database driver JAR file, contact your database vendor. Database driver file should be installed to the `<pf_install>/pingfederate/server/default/lib` directory. You must restart the server after installing the driver.

1. Go to the **Server Configuration > Data Stores** screen.

Task	Steps
Configure a new data store	<ol style="list-style-type: none"> 1. On the Manage Data Stores screen, click Add New Data Store. 2. On the Data Store Type screen, select Database.
Modify an existing data store	<ol style="list-style-type: none"> 1. On the Manage Data Stores screen, click the applicable data store by its name. 2. On the Summary screen, click Database Config.

2. On the **Database Config** screen, configure your JDBC connection, as described in the following table.

Field	Description
JDBC URL (Required)	<p>The location of the JDBC database. The structure of the JDBC URL varies depending on the vendor; for example:</p> <ul style="list-style-type: none"> • <code>jdbc:mysql://databaseservername/databasename</code> • <code>jdbc:oracle:thin:@databaseservername:databasename</code> • <code>jdbc:postgresql://databaseservername/databasename</code>

Field	Description
	<ul style="list-style-type: none"> <code>jdbc:sqlserver://databaseservername;databaseName=databasename</code> <p>where <i>databaseservername</i> is the DNS host name (or IP) of the server hosting the database, and <i>databasename</i> is the name of a database on that server.</p> <p> 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:</p> <pre>?autoReconnect=true</pre>
Driver Class (Required)	<p>The name of the driver class used to communicate with the source database; for example:</p> <ul style="list-style-type: none"> <code>com.microsoft.sqlserver.jdbc.SQLServerDriver</code> <code>com.mysql.jdbc.Driver</code> <code>oracle.jdbc.OracleDriver</code> <code>org.postgresql.Driver</code> <p>The driver class name should be supplied by the database software vendor in a JAR file.</p>
Username (Required)	The name that identifies the user when connecting to the database.
Password	The password needed to access the database.
Validate Connection SQL (Optional, but recommended)	<p>A simple SQL statement used by PingFederate® at runtime to verify that the database connection is still active and to reconnect if needed.</p> <p>If a SQL statement is not provided here, PingFederate may not be able to reconnect to the database if the connection is broken.</p> <p> Important: Ensure that the SQL statement is valid for your database; for example:</p> <ul style="list-style-type: none"> <code>SELECT 1 from dual</code> (for Oracle Database or Oracle MySQL) <code>SELECT getdate()</code> (for Microsoft SQL Server) <code>SELECT 1</code> (for PostgreSQL) <p> Tip: To use this feature for Oracle MySQL, you must also add the <code>?autoReconnect=true</code> query parameter to the JDBC URL.</p>
Mask Values in Log	Determines whether all attribute values returned from this data store should be masked in PingFederate log files.
Allow Multi-Value Attributes	When selected (the default), indicates that this JDBC data store can select more than one record from a column and return the results as a multi-value attribute. Otherwise, a query returns only the first value in the column.

3. Optional: Click **Advanced** to configure additional settings.

- On the **Advanced Database Options** screen, click **Apply Defaults** to view or restore default values.
- Configure advanced settings, as described in the following table.

Field	Description
Minimum Pool Size	The smallest number of database connections that can remain in the pool for the given data store.

Field	Description
	Note that PingFederate does not establish the connection pool for the given data store until it receives a request that requires one or more attributes from that data store.
Maximum Pool Size	The largest number of database connections that can remain in the pool for the given data store.
Blocking Timeout (ms)	The amount of time a request waits to get a connection from the connection pool before it fails.
Idle Timeout (min)	The length of time the connection can be idle in the pool before it is closed. Note that PingFederate maintains the minimum connection pool for the given data store once the pool is established.

4. Click **Save** to retain the configuration.




Configure an LDAP connection




You configure access to user attributes on your LDAP server by providing basic LDAP information.

1. Go to the **Server Configuration > Data Stores** screen.

Task	Steps
Configure a new data store	<ol style="list-style-type: none"> 1. On the Manage Data Stores screen, click Add New Data Store. 2. On the Data Store Type screen, select LDAP.
Modify an existing data store	<ol style="list-style-type: none"> 1. On the Manage Data Stores screen, click the applicable data store by its name. 2. On the Summary screen, click LDAP Configuration.

2. On the **LDAP Configuration** screen, configure your JDBC connection, as described in the following table.

Field	Description
Hostname(s) (Required)	<p>The DNS name or IP address of the data store, which may include a port number; example: 181.20.42.130:389.</p> <p>For failover, you can enter multiple LDAP servers, each separated by a space.</p> <p> Note: If multiple LDAP servers are specified, each server must be accessible using the same user DN and password (unless the Bind Anonymously check box is selected).</p>
LDAP Type (Required)	<p>If you are using this data store for outbound provisioning and your LDAP store is PingDirectory™, Microsoft Active Directory (AD), or Oracle Directory Server, select the applicable type from the list, such that PingFederate® can pre-configure many provisioning settings automatically.</p> <p> Tip: If you are using outbound provisioning and your LDAP server is <i>not</i> one of the three aforementioned LDAP servers, you may define a custom LDAP Type to streamline the provisioning configuration on the Outbound Provisioning > Channel > Source Settings screen.</p> <p>The LDAP type is also used to enable password-change messaging between AD and PingFederate when an HTML Form Adapter instance is used.</p>
Bind Anonymously	<p>Select this check box if your LDAP directory server supports anonymous binding and if no credentials are needed to access the data store.</p> <p> Tip: If you choose an anonymous binding, ensure that this access level provides permission to search the directory for user-account information.</p>

Field	Description
	<p>For inbound provisioning, because PingFederate needs to manage local user records, your LDAP directory server likely requires a specific service account to handle the communication between PingFederate and the target data store.</p> <p>When selected, user DN and password are not required.</p>
User DN	<p>The username credential required to access the data store.</p> <p> Important: The service account must have permission to search the directory for user-account information. If your use cases involve reading from the LDAP server without creating, updating, or deleting any records, consider using a service account with read-only access.</p> <p>For inbound provisioning, a service account with permission to create, read, update, and delete (if applicable) users (and groups if applicable) is required.</p> <p>When connecting to an AD LDAP server running on Microsoft Windows Server 2008 R2 (or newer), enter an AD User account; do not use a Computer account.</p> <p>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 data store. For more information, see Configure proxied authorization on page 162.</p>
Password	The password credential required to access the data store.
Use LDAPS	<p>Connects to the LDAP data store using LDAPS. This selection applies equally to any secondary servers specified in the Hostname(s) field.</p> <p> Important: We recommend that all LDAP connections be secured by using LDAPS.</p> <p> Note: If you want to enable the password changes, password reset, or account unlock features in the HTML Form Adapter against an AD LDAP server, you <i>must</i> secure the data store connection to your AD LDAP server using LDAPS; AD requires this level of security to allow password changes.</p>
Mask Values in Log (Checkbox)	Determines whether all attribute values returned from this data store should be masked in PingFederate log files.

3. Optional: Click **Advanced** to configure additional settings in the **Advanced LDAP Options** screen.

4. Click **Save** to retain the configuration.

Configure 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 data store. By doing so, users cannot reset password if their account has been disabled or if they have not been granted the permission to change their password.

Please refer to the following resources to configure proxy authentication for the LDAP service account.

- For PingDirectory, see the [PingDirectory Server Administration Guide](#) and search for *Proxied Authorization*.
- For Oracle Directory Server, see the [Oracle Fusion Middleware Deployment Planning Guide](#) (docs.oracle.com/cd/E29127_01/doc.111170/e28974/security-requirements.htm#aallhm) and search for *Proxy Authorization*.

Note that Microsoft Active Directory does not support proxied authorization (see the [Microsoft Active Directory Technical Specification](https://msdn.microsoft.com/library/cc223358.aspx) at msdn.microsoft.com/library/cc223358.aspx).

For general information about proxied authorization, please refer to [RFC4370](#).

Set advanced LDAP options

PingFederate® maintains a search pool and a bind pool for each LDAP data store instance for optimal performance. The search pool is meant for LDAP directory searches. The bind pool is meant for LDAP bind authentication purposes.

Use the **Advanced LDAP Options** screen to change default pool settings for the related data store. These settings are applicable to both the search pool and the bind pool.

1. Optional: Click **Apply Defaults** to view or restore default values.



Tip: The default values are conservative based on the default thread-pooling limits allowed by the installed PingFederate web container. (These limits are under "Server Thread Pool" in the `jetty-runtime.xml` file, located in the `<pf_install>/pingfederate/etc` directory.)

If any changes are made to thread pooling, we recommend updating settings as outlined in the following table.

2. Configure advanced settings, as described in the following table.

Field	Description
Test Connection on Borrow	Indicates whether objects are validated before being borrowed from the pool.
Test Connection on Return	Indicates whether objects are validated before being returned to the pool.
Create New Connection If Necessary	Indicates whether temporary connections can be created 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 data store instance may fail.
Verify LDAPS Hostname	Indicates whether to verify the hostname of the LDAP server matches the Subject (CN) or one of the Subject Alternative Names from the certificate. Important: We recommend to verify LDAPS hostname for all LDAPS connections.
Minimum Connections (Required)	The smallest number of connections that can remain in each pool. A minimum value of 1 creates two connections: one connection in the search pool and one connection in the bind pool. Note: For optimal performance, the value for this setting should be equal to 50% of the <code>maxThreads</code> value in the Jetty server configuration (see the Tip in step 1). Note that PingFederate does not establish the connection pool for the given data store until it receives a request that requires one or more attributes from that data store.
Maximum Connections (Required)	The largest number of active connections that can remain in each pool (not including the temporary connections that are managed automatically when the Create New Connection If Necessary check box is selected). The value must be greater than or equal to the Minimum Connections value. Note: For optimal performance, the value for this setting should be equal to 75% to 100% of <code>maxThreads</code> value in the Jetty server configuration (see the Tip in step 1).

Field	Description
Maximum Wait (Milli) (Required)	The maximum number of milliseconds the pool waits for a connection to become available when trying to obtain a connection from 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).
Time Between Eviction (Milli) (Required)	The frequency in milliseconds that the evictor cleans up the connections in the pool. A value of -1 disables the evictor.
Read Timeout (Milli) (Required)	The maximum number of milliseconds a connection waits for a response to be returned before producing an error. A value of -1 causes the connection to wait indefinitely.
Connection Timeout (Milli) (Required)	The maximum number of milliseconds that a connection attempt should be allowed to continue before returning an error. A value of -1 causes the pool to wait indefinitely.

- Optional: Click **Next** to specify LDAP binary attributes in the **LDAP Binary Attributes** screen.
- Click **Save** to retain the configuration.

Specify LDAP binary attributes

PingFederate[®] allows you to specify attributes that must be handled as binary data for use in attribute contract fulfillment.



Note: Binary data cannot be used in an assertion. Encoding must be applied and is handled on a connection basis. When binary attributes are selected for attribute mapping, the administrative console prompts you to select an encoding type for each binary attribute.

- On the **LDAP Binary Attributes** screen, add, update, or remove binary attributes.
- Click **Save** to retain the configuration.

Define a custom LDAP type for outbound provisioning

If you are using outbound provisioning and your user-management LDAP server is not PingDirectory[™], Microsoft Active Directory, or Oracle Directory Server, you can define a custom LDAP type for PingFederate[®] to use to streamline the provisioning configuration.

When the LDAP server is defined, its type appears in the **LDAP Type** list on the **LDAP Configuration** screen.

- Copy and rename the `sample.template.txt` file located in the `<pf_install>/pingfederate/server/default/conf/template/ldap-templates` directory.
- Change the `template.name` property value in the new template file.

The `template.name` property value appears in the **LDAP Type** list on the **LDAP Configuration** screen when you save the template.

- Modify other property values in the file to match the corresponding configuration of your LDAP server.

These properties correspond to the fields shown on the **Outbound Provisioning > Channel > Source Settings** screen, where

These are the properties that the provisioner uses to determine when user information is added, changed, or removed. from the outbound provisioning setup, specifically the **Source Settings** screen. The provisioner uses these fields to determine when user information is added, changed, or removed.

- 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.

- Go to the **Server Configuration > Data Store** screen.

You may create a new LDAP connection to your LDAP directory server by using the newly defined custom LDAP type. (Select the custom LDAP type from the list in the **Data Store Type** screen.)

Configure a custom data store

Developers can use the PingFederate® Custom Source SDK to create specific drivers for non-JDBC/LDAP data stores (or more sophisticated JDBC/LDAP queries) including, for example, flat files or SOAP-connected databases. Furthermore, custom data stores may be written to perform configuration assistance or validation actions, such as testing a connection to a database. Actions may also include generation of parameters that might need to be set manually in a configuration file.

Once the data-store driver (JAR) file is installed, you can select it on the **Custom Data Store Type** screen.

1. Implement your custom data store, copy the data-store driver (JAR) file to the `<pf_install>/pingfederate/server/default/deploy` directory, and then restart the PingFederate service.
2. Go to the **Server Configuration > Data Stores** screen.

Task	Steps
Configure a new data store	<ol style="list-style-type: none"> 1. On the Manage Data Stores screen, click Add New Data Store. 2. On the Data Store Type screen, select Custom.
Modify an existing data store	<ol style="list-style-type: none"> 1. On the Manage Data Stores screen, click the applicable data store by its name. 2. On the Summary screen, click the applicable screen to modify the settings.

3. On the **Customer Data Store Type** screen, configure the instance of your custom data store as follows:
 - a) Enter a unique name in the **Data Store Instance Name** field.

You can create more than one instance of the same custom data store for use with different partners, as needed.
 - b) Select the desired custom data store type from the list.

If the list is empty, verify that the data-store driver is implemented and installed successfully.
 - c) Select the **Mask Values in Log** check box if all attribute values returned from this data store should be masked in PingFederate log files.

For example, after building and deploying the sample from the `<pf_instal>/pingfederate/sdk/plugin-src/custom-data-store-example` directory, you can create an instance of the **Sample SDK Properties Data Store**:

Manage Data Stores | Data Store

Data Store Type	Custom Data Store Type	Configure Attribute Source Adapter Instance	Summary
-----------------	------------------------	---	---------

Enter a descriptive name, a system-wide unique ID, and select the custom data store adapter to use.

DATA STORE INSTANCE NAME	<input type="text" value="My first custom data store"/>
DATA STORE TYPE	<input type="text" value="Sample SDK Properties Data Store"/>
<input type="checkbox"/>	MASK VALUES IN LOG

Cancel

Previous

4. On the **Configure Attribute Source Adapter Instance** screen, configure the instance of your custom data store.

This screen varies depending on the implementation. For example, the following screenshot is based on the sample from the `<pf_install>/pingfederate/sdk/plugin-src/custom-data-store-example` directory.

Manage Data Stores | Data Store

Data Store Type	Custom Data Store Type	Configure Attribute Source Adapter Instance	Summary
-----------------	------------------------	---	---------

Configure the Custom Source Adapter.

Configuration settings for the sample properties data store.

Field Name	Field Value	Description
PATH TO PROPERTIES DIRECTORY	<input type="text"/>	The path specifies which directory the properties files are located. Each properties file in the directory should contain entries for 'favoriteMovie', 'favoriteBook' and 'favoriteSong'.

5. Click the related link on the **Actions** screen to invoke an action (when applicable).
6. Click **Save** to retain the configuration.

Define an account-linking data store

When an SP is configured to use account linking for an IdP connection, PingFederate uses an embedded HyperSQL database as the account-link repository (see [Account linking](#) on page 69). This default implementation does not require any changes to PingFederate to support account linking. However, you can manually customize PingFederate to store account links in a different data store—either a different database or an LDAP directory. You might want to do this for any of several reasons, including:

- You are running a cluster of PingFederate runtime engines (see the PingFederate [Server clustering guide](#) on page 533). This scenario requires that you use an external database or directory for account links to ensure proper local user lookup.
- You have performance or scalability requirements that exceed the HyperSQL database's capabilities.
- You and your federation partner previously established a different system for creating and mapping opaque pseudonyms, and PingFederate needs access to the system.

Change the account-linking database

Changing the default database involves creating a table in your JDBC database to support account linking, and modifying PingFederate configuration XML files to use the database.

To create a database table for account linking:

- Run one of the table-setup scripts provided in the directory:
`<pf_install>/pingfederate/server/default/conf/account-linking/sql-scripts`

If a script is not provided for your database, you can derive the setup from information available in any of the other scripts.

To change the account-linking database:

1. If you have not already done so, create a connection to the database you want to use (see [Configure a JDBC database connection](#) on page 159).

Be sure to save the configuration on the Manage Data Stores screen.

2. Any time after saving the database connection, return to the Manage Data Stores screen from the Main Menu. (To reach the screen, click **Data Stores** under System Settings.)
3. Copy the System ID for the database you want.
4. In the directory `<pf_install>/pingfederate/server/default/data/config-store`, open the file:


```
org.sourceid.saml20.service.impl.AccountLinkingServiceDBImpl.xml
```

Between the XML tags for the item named `PingFederateDSJNDIName`, insert the System ID you copied at [Step 3](#) and save the file.

5. Start or restart PingFederate.



Note: If you are running PingFederate in a cluster, push the new configuration to other server nodes. For more information, see the PingFederate [Server clustering guide on page 533](#).

Change the default data store to use LDAP

Changing the default data store to use LDAP involves modifying PingFederate configuration XML files to use the LDAP directory.

To use an LDAP directory:

1. If you have not already done so, create a connection to the LDAP data store you want to use (see [Configure an LDAP connection](#) on page 161).

Be sure to save the configuration on the Manage Data Stores screen.

2. Any time after saving the LDAP data store connection, return to the Manage Data Stores screen from the Main Menu. (To reach the screen, click **Data Stores** under System Settings.)
3. Copy the System ID for the data store you want.
4. In the directory `<pf_install>/pingfederate/server/default/conf/META-INF`, open the file: `hivemodule.xml`

Locate the Service-Point ID for `AccountLinkingService` and change the value of the `create-instance` class to:

```
org.sourceid.saml20.service.impl.AccountLinkingServiceLDAPImpl
```

5. In the directory `<pf_install>/pingfederate/server/default/data/config-store`, open the file: `org.sourceid.saml20.service.impl.AccountLinkingServiceLDAPImpl.xml`


Insert the following values between the XML tags for the these items:


- `PingFederateDSJNDIName`: System ID you copied at [Step 3](#)
- `UserSearchBase`: LDAP location where searches begin—for example, `CN=Users,DC=LDAPDir,DC=com`
- `UsernameAttribute`: LDAP attribute that represents the user identifier—for example, Active Directory is `sAMAccountName`
- `AccountLinkDataAttribute`: LDAP attribute used to store account linking data



Note: The `AccountLinkDataAttribute` can be any multi-valued string attribute on a user object class. We recommend that you extend the LDAP schema with a custom attribute for use here. See this [article](http://msdn.microsoft.com/library/ms676900(v=VS.85).aspx) from Microsoft for further information on extending the Active Directory schema ([msdn.microsoft.com/library/ms676900\(v=VS.85\).aspx](http://msdn.microsoft.com/library/ms676900(v=VS.85).aspx)).

6. Start or restart PingFederate.
7. If you are running PingFederate in a cluster, push the new configuration to other server nodes (see the PingFederate [Server clustering guide on page 533](#))

 **Important:** You must manually apply the changes made in [Step 4](#) to the `hivemodule.xml` file on each server node in a cluster and then start or restart PingFederate.

 **Note:** User accounts to be linked must exist in the LDAP directory prior to establishing the account link. The Account Linking service does not add users to the LDAP data store but simply updates `AccountLinkDataAttribute` for a given user.

Define an OAuth grant data store

PingFederate® uses its internal HyperSQL database by default to maintain persistent grants (if any) for the OAuth AS.

Persistent grants can result from the following OAuth use cases:


- OAuth clients using the **Refresh Token** grant type in conjunction with either the **Authorization Code** or **Resource Owner Credentials** grant type.

If the use cases involve mapping attributes from authentication sources (IdP adapter instances or IdP connections) or Password Credential Validator (PCV) instances to the access tokens (directly or through persistent grant extended attributes), such attributes and their values are stored along with the persistent grants so that they can be reused when clients subsequently present refresh tokens for new access tokens.

- OAuth clients using the **Implicit** grant type *and* the **Reuse Existing Persistent Access Grants for Grant Types** check box is selected in the **OAuth Settings > Authorization Server Settings** screen.

If the use cases involve mapping attributes from authentication sources or PCV instances to the access tokens (directly or through persistent grant extended attributes), attribute values are obtained at runtime for each token request. No attributes or their values are stored with the persistent grants.

Persistent grants (and the associated attributes and their values, if any) remain valid until the grants expired or are explicitly revoked.


 **Important:** When persistent grants are expected, administrators should consider changing this configuration to use an external secured database or an LDAP directory server for production standalone deployments. For server clustering, an external database or an LDAP directory server is required, because the HyperSQL database cannot be shared across other PingFederate engine nodes.

Changing the default storage for OAuth persistent grants involves creating the required data structure on the external database server or LDAP directory server and modifying a couple PingFederate configuration XML files.

PingFederate also supports other storage solutions through the PingFederate SDK. The process involves implementing the `AccessGrantManager` interface and modifying a PingFederate configuration XML file to point to the custom class.


Configure an external database for grant storage

1. Run the table-setup scripts for your database server provided in the `<pf_install>/pingfederate/server/default/conf/access-grant/sql-scripts` directory.

 **Note:** If scripts are not provided for your database server, you can derive the setup from information available in any of the other scripts.

2. If you have not already done so, create a connection to the database you want to use (see [Configure a JDBC database connection](#) on page 159).
3. Go to the **Server Configuration > Data Stores** screen.
4. Copy the **System ID** value for the database you want.
5. Edit the `org.sourceid.oauth20.token.AccessGrantManagerJdbcImpl.xml` file in the `<pf_install>/pingfederate/server/default/data/config-store` directory.

Change the value of the item named `PingFederateDSJNDIName` to the **System ID** value you copied, and then save the file.


 **Note:** For a clustered PingFederate® environment, edit the `org.sourceid.oauth20.token.AccessGrantManagerJdbcImpl.xml` file on the console node.

6. Edit the `hivemodule.xml` file in the `<pf_install>/pingfederate/server/default/conf/META-INF` directory.

- a) Locate the `AccessGrantManager` block:

```
<!-- Service for storage of access grants -->
<service-point id="AccessGrantManager"
  interface="com.pingidentity.sdk.accessgrant.AccessGrantManager">
  <create-instance
    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.

 **Important:** For a clustered PingFederate environment, edit the `hivemodule.xml` file on each node.

7. Start or restart PingFederate.

For a clustered PingFederate environment, click **Replicate Configuration** on the **Server Configuration > Cluster Management** screen to push the modified `org.sourceid.oauth20.token.AccessGrantManagerJdbcImpl.xml` file to all engine nodes, and then restart each engine node to activate the change.

PingFederate automatically removes expired grants once a day. To fine-tune the frequency and the number of expired grants to be removed, see [Manage expired OAuth persistent grants](#) on page 140.


Configure an LDAP directory for grant storage

Specific schema objects are required in order for PingFederate to store grants on your LDAP directory server. LDIF scripts are provided for supported LDAP directory servers (see [Deployment options](#) on page 13).

1. Review the LDIF scripts for your LDAP 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 LDAP directory server.
3. Run the LDIF scripts to update your LDAP schema.
4. If you have not already done so, create a connection to your LDAP directory server (see [Configure an LDAP connection](#) on page 161).
5. Go to the **Server Configuration > Data Stores** screen.
6. Copy the **System ID** value for the LDAP data store connecting to your LDAP directory server.
7. Edit a configuration file in the `<pf_install>/pingfederate/server/default/data/config-store` directory.

LDAP directory server	Configuration file
PingDirectory™	<code>org.sourceid.oauth20.token.AccessGrantManagerLDAPPingDirecto</code>
Microsoft Active Directory	<code>org.sourceid.oauth20.token.AccessGrantManagerLDAPADImpl.xml</code>
Oracle Directory Server Enterprise Edition	<code>org.sourceid.oauth20.token.AccessGrantManagerLDAPOracleImpl.</code>

- a) Use the **System ID** value you copied as the value of the item named `PingFederatedSjndiName`.
- b) Enter a value for the item named `SearchBase`.

 **Tip:** This is the LDAP search base that points to the `accessGrants` location. For more information, see 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/ldif-scripts` directory.

d) Save the file.

 **Note:** For a clustered PingFederate environment, edit the configuration file on the console node.

8. Edit the `hivemodule.xml` file in the `<pf_install>/pingfederate/server/default/conf/META-INF` directory.


a) Locate the `AccessGrantManager` block:

```
<!-- Service for storage of access grants -->
<service-point id="AccessGrantManager"
  interface="com.pingidentity.sdk.accessgrant.AccessGrantManager">
  <create-instance
    class="org.sourceid.oauth20.token.AccessGrantManagerJdbcImpl"/>
</service-point>
```

b) Set the value of the class attribute to one of the following values:

LDAP directory server	Class value
PingDirectory	<code>org.sourceid.oauth20.token.AccessGrantManagerLDAPPingDirect</code>
Microsoft Active Directory	<code>org.sourceid.oauth20.token.AccessGrantManagerLDAPADImpl</code>
Oracle Directory Server Enterprise Edition	<code>org.sourceid.oauth20.token.AccessGrantManagerLDAPOracleImpl</code>

c) Save the file.

 **Important:** For a clustered PingFederate environment, edit the `hivemodule.xml` file on each node.

9. Start or restart PingFederate.

For a clustered PingFederate environment, click **Replicate Configuration** on the **Server Configuration > Cluster Management** screen to push the modified configuration file to all engine nodes, and then restart each engine node to activate the changes.

Improve performance with indexes

If you use PingDirectory™ to store OAuth persistent grants, the following attributes should be indexed to ensure that access grant queries perform efficiently.

Attribute name	Index type
<code>accessGrantGuid</code>	equality
<code>accessGrantUniqueUserIdentifier</code>	equality
<code>accessGrantHashedRefreshTokenValue</code>	equality
<code>accessGrantClientId</code>	equality
<code>accessGrantExpires</code>	ordering

These indexes can be created using the PingDirectory's `dsconfig` utility. The `dsconfig` utility is interactive. You can also provide inputs as command arguments.

For example, the following samples create the `accessGrantGuid` and `accessGrantExpires` 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 accessGrantExpires \
```

```
--set index-type:ordering
```

After adding the indexes, use the `rebuild-index` utility to build each of the five indexes. For instance, the following sample builds the `accessGrantGuid` index.


```
$ bin/rebuild-index \
  --baseDN "dc=example,dc=com" \
  --index accessGrantGuid
```

For more information, see the "Managing Local DB Indexes" section in the [PingDirectory Server Administration Guide](#).

Use a custom solution for grant storage

1. Implement the `AccessGrantManager` interface.

For more information, refer to the Javadoc for the `AccessGrantManager` interface, the `SampleAccessGrant.java` sample implementation, and the [SDK developer's guide](#) for build and deployment information.

 **Tip:** The Javadoc for `PingFederate` and the sample implementation are located under the `<pf_install>/pingfederate/sdk` directory.


2. Edit the `hivemodule.xml` file in the `<pf_install>/pingfederate/server/default/conf/META-INF` directory.

a) Locate the `AccessGrantManager` block:

```
<!-- Service for storage of access grants -->
<service-point id="AccessGrantManager"
  interface="com.pingidentity.sdk.accessgrant.AccessGrantManager">
  <create-instance
    class="org.sourceid.oauth20.token.AccessGrantManagerJdbcImpl"/>
</service-point>
```

b) Set the value of the class attribute to the name of your class.

c) Save the file.

 **Important:** For a clustered `PingFederate` environment, edit the `hivemodule.xml` file on each node.


3. Deploy the required program files of your custom implementation to all `PingFederate` servers.

4. Start or restart `PingFederate`.

Define an OAuth client data store

By default, `PingFederate`[®] stores OAuth clients entered into the system through the administrative console or the REST-based application programming interface in XML files.

Alternatively, you can configure `PingFederate` to use a database for OAuth client records, allowing the records to be managed via the administrative console, the REST-based application programming interface, or an OAuth Client Web Service. When database usage is enabled, `PingFederate` uses its internal HyperSQL database by default. Administrators should consider changing this configuration to use an external secured database for production standalone deployments.

 **Important:** For server clustering, an external database *is* required, because the HyperSQL database cannot be shared across multiple `PingFederate` engine nodes.

Lastly, you may store OAuth clients in other storage medium by implementing the `ClientStorageManager` interface, available through the `PingFederate` SDK.

Enable database storage for OAuth clients


1. Edit the `<pf_install>/pingfederate/server/default/conf/META-INF/hivemodule.xml` file.

2. Locate the ClientManager service point:

```
<!-- Service for storing OAuth client configuration. -->
<service-point id="ClientManager"
  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.ClientManagerGenericImpl ...
    -->
    <construct
      class="org.sourceid.oauth20.domain.ClientManagerJdbcImpl"/>
  </invoke-factory>
</service-point>
```

Update the class attribute value to `org.sourceid.oauth20.domain.ClientManagerJdbcImpl`.

If you subsequently want to switch back to the default XML file storage implementation, revert the class attribute value to `org.sourceid.oauth20.domain.ClientManagerXmlFileImpl`.

 **Caution:** If clients are already defined using the default XML implementation, the records are not moved to a database: they must be recreated. Likewise, if the configuration is changed back to the default, database records are not ported to XML files.

3. Restart PingFederate.

4. If you are running PingFederate in a server cluster, repeat the previous steps for each server node.

For server clustering, an external database *is* required, because the HyperSQL database cannot be shared across multiple PingFederate engine nodes. For production standalone deployments, administrators should consider using an external secured database as well.

Follow the remaining steps to set up an external database as the storage medium for your OAuth clients.

5. For the database server that you want to use, run the associated table-setup scripts provided in the `<pf_install>/pingfederate/server/default/conf/oauth-client-management/sql-scripts` directory.
6. If you have not already done so, create a JDBC data store to your database server on the **Server Configuration > Data Stores** screen (see [Configure a JDBC database connection](#) on page 159).
7. Copy the system ID of the applicable JDBC data store from the **Server Configuration > Data Stores** screen.
8. Edit the `<pf_install>/pingfederate/server/default/data/config-store/org.sourceid.oauth20.domain.ClientManagerJdbcImpl.xml` file and replace the value of the `c:item` element with the system ID from the previous step.

For example, if the system ID from the previous step is

`JDBC-123456789ABCDEF123456789ABCDEF123456ABCD`, update the

`org.sourceid.oauth20.domain.ClientManagerJdbcImpl.xml` file as follows:

From:


```
<?xml version="1.0" encoding="UTF-8"?>
<c:config xmlns:c="http://www.sourceid.org/2004/05/config">
  <c:item name="PingFederatedSJNDIName">PFDefaultDS</c:item>
</c:config>
```

(PFDefaultDS is the default value.)

To:

```
<?xml version="1.0" encoding="UTF-8"?>
<c:config xmlns:c="http://www.sourceid.org/2004/05/config">
```

```
<c:item
  name="PingFederateDSJNDIName">JDBC-123456789ABCDEF123456789ABCDEF123456ABCD</
c:item>
</c:config>
```

 **Note:** If you are running PingFederate in a server cluster, make this modification on the console node.

9. Restart PingFederate.


10. If you are running PingFederate in a server cluster, replicate this new configuration to other engine nodes on the **Server Configuration > Cluster Management** screen.

Enable custom storage for OAuth clients

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.

Refer to the PingFederate Javadoc for information about the `ClientStorageManagerV2` interface, the `SampleClientStorage.java` file for a sample implementation, and the [SDK developer's guide](#) for build and deployment information.

 **Tip:** The PingFederate Javadoc and the sample implementation are located underneath the `<pf_install>/pingfederate/sdk` directory.

2. Edit the `<pf_install>/pingfederate/server/default/conf/META-INF/hivemodule.xml` file.

3. Locate the `ClientStorageManager` service point:

```
<!-- Service for using custom OAuth client storage implementation. -->
<service-point id="ClientStorageManager"
  interface="com.pingidentity.sdk.oauth20.ClientStorageManagerV2">
  <invoke-factory>
    <!--
      Provide the class name for the custom OAuth client storage
      implementation.
    -->
    <construct
      class="org.sourceid.oauth20.domain.NoOpClientStorageManager"/>
  </invoke-factory>
</service-point>
```

Update the class attribute value with the name of the class implementing the `ClientStorageManagerV2` interface.

4. Locate the `ClientManager` service point:

```
<!-- Service for storing OAuth client configuration. -->
<service-point id="ClientManager"
  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.ClientManagerGenericImpl ...
  -->
  <construct
    class="org.sourceid.oauth20.domain.ClientManagerXmlFileImpl"/>
  </invoke-factory>
```



```
</service-point>
```

Update the class attribute value to `org.sourceid.oauth20.domain.ClientManagerGenericImpl`.

If you subsequently want to switch back to the default XML file storage implementation, revert the class attribute value to `org.sourceid.oauth20.domain.ClientManagerXmlFileImpl`.



Caution: If clients are already defined using the default XML implementation, the records are not moved to the custom storage: they must be recreated. Likewise, if the configuration is changed back to the default, OAuth clients stored in the custom storage are not ported to XML files.

5. Restart PingFederate.
6. If you are running PingFederate in a server cluster, repeat the previous steps for each server node.

IdP discovery

IdP discovery refers to the process of identifying an end-user's IdP dynamically during an SP-initiated SSO event.

An SP can 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.

PingFederate® also provides two kinds of IdP discovery:

- SAML 2.0 standard IdP Discovery
- Proprietary IdP discovery using a persistent cookie written by an SP PingFederate server

Configure IdP discovery using a persistent cookie

PingFederate®'s proprietary IdP-discovery method makes use of an IdP persistent reference cookie (IPRC) to track the identity provider with whom a user last authenticated. There are three significant differences between standard IdP Discovery and the IPRC method:

- Standard IdP Discovery may be used only with SAML 2.0; the IPRC may be used with any federation protocol.
- The common domain cookie (CDC) may 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 readable by both federation partners; the IPRC is set by the SP, using information in the SAML assertion, and cannot be accessed by the IdP.

Note that the deployed connection configuration between SP and IdP partners must include SP-initiated SSO.

1. Edit the `org.sourceid.websso.profiles.sp.IdpIdCookieSupport.xml` file located in the `<pf_install>/pingfederate/server/default/data/config-store` directory.
2. Set the value of `EnableIdpIdCookie` to `true`.
3. Optional: Modify the remaining elements in the configuration, as described in the following table:

Field	Description
<code>IdpIdCookieName</code>	The name of the IPRC set by the SP installation (default: <code>IdPIId</code>). Note that the cookie name cannot contain any of the following characters: <code>&</code> , <code>></code> , <code><</code> , comma, semicolon, space.
<code>IdpIdCookieLifeTimeInDays</code>	The lifetime for the cookie (default: 365 days, maximum: 24,855 days). The browser will delete the cookie when the period is expired.
<code>ShowIdpSelectionList</code>	If set to <code>true</code> (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 <code>false</code> , the SP installation generates an error page.

4. Start or restart PingFederate.



Note: 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.

Configure standard IdP Discovery

SAML 2.0 IdP Discovery provides a cookie-based look-up mechanism used to identify a user's IdP dynamically during an SP-initiated SSO event, when the IdP is not otherwise specified. This mechanism can be helpful, in particular, in cases where an SP might be a hub for several IdPs in an identity federation.

When a user requests access to a protected resource on the SP, common-domain 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
- Common domain cookie reader

Each of these roles is necessary to support IdP Discovery. The roles may 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.

1. If you have not done so, go to the **Server Configuration > Server Settings > Roles & Protocols** screen and select the IdP Discovery role.
2. Go to the **Server Configuration > IdP Discovery** screen, and then click **Configure IdP Discovery**.
3. On the **Domain Cookie Settings** screen, choose the discovery role or roles that PingFederate will play.

The choices that appear on this screen depend on whether PingFederate is acting as an SP, an IdP, or both; or as an IdP Discovery server only in the **Server Configuration > Server Settings > Roles & Protocols** screen.

4. On the **Common Domain Service** screen, configure as follows:

Field	Description
Base URL	Enter the base URL of the 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	Enter and confirm the pass phrase that web applications must use to access the domain.

5. On the **Local Common Domain Server** screen, 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 (.); for example: <code>.example.com</code>
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	Enter and confirm the pass phrase that web applications must use to access the domain.

6. On the **Summary** screen, review and modify settings as needed, and then save your configuration.
7. Perform one of the following actions to enable the setting of the common domain cookie at runtime:
 - Ensure that, prior to launching any SSO events, the web application that implements IdP Discovery sets the cookie using the `/idp/writetcdc.ping` application endpoint intended for that purpose.
 - Enable setting the cookie at runtime during SSO events by selecting the **IdP Discovery** check box in the **Connection Options** screen for the desired SP connection.

Configure redirect validation

Several SP adapters can be configured 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 open the possibility that a third party (with specific knowledge of the IdP, the SP, or both; PingFederate endpoints; and PingFederate configuration) might be able to obtain and use valid security tokens to gain improper access to the target resource.

This potential security threat would involve using well-formed SSO or SLO links to start an SSO or SLO request for a resource at the SP site. However, the target resource designated in the link would be intended to intercept the security token by redirection to a malicious website.

To prevent such an attack, PingFederate provides a means of validating SSO and SLO transactions to ensure that the designated target resource exists in a domain controlled by the SP through a list of configurable URLs.



Note: Validating target resource for SSO is applicable for IdP connections, adapter-to-adapter mappings and SAML 2.0 IdP Discovery.

The following default target URLs are always allowed:

- The default target URL for any IdP connections (see [Configure default target URLs](#) on page 363)
- The default target URL for any adapter-to-adapter mappings (see [Configure a default target URL \(optional\)](#) on page 422)
- The SP default URL for successful SSO (see [Configure default URLs](#) on page 336)
- The IdP default URL for successful SLO (see [Configure a default URL and error message](#) on page 280)

They do not need to be entered into the list manually.



Tip: PingFederate is also capable of validating the error resource (`InErrorResource`) parameter for the same reason. For more information, see [IdP endpoints](#) on page 448, [SP endpoints](#) on page 451 and [System-services endpoints](#) on page 460.



Important: PingFederate enables target resource validation for SSO and SLO, as well as error resource URLs, and requires HTTPS in new installations by default.

For backward compatibility, PingFederate Upgrade Utility does not enable these options if they were not selected in the previous PingFederate installation. Although optional, it is strongly recommended to enable target and error resource validation (including the HTTPS option) and to enter all expected resources to prevent unauthorized access.

To reach this screen:

1. Click **Server Configuration** on the Main Menu.
2. Click **Redirect Validation** under System Settings.

To enable target resource validation for SSO and/or SLO:


- Select the SSO and/or SLO check box(es).

To enable error resource validation:

- Select the Enable InErrorResource validation check box.

To enter expected resources:

1. Indicate whether to 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.

2. Enter the domain or IP address containing the expected target resource under **Valid Domain Name**.

Use the domain only without qualifiers; for example: `example.com`


Using an initial wildcard and period for a domain name will cover multiple sub domains; for example:

`*.example.com` covers `hr.example.com` or `email.example.com` but not `example.com`.

 **Important:** While using an initial wildcard provides the convenience of allowing multiple sub domains with one entry, consider adding individual sub domains to limit the redirection to a list of known hosts.


(Optional) Enter the exact path of the resource (case-sensitive) under **Valid Path**. Starts 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. For example:

`/inbound/Consumer.jsp` allows `/inbound/Consumer.jsp` but rejects `/inbound/consumer.jsp`

 **Tip:** You can also enter multiple query parameters with or without a fragment.

For example: `/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`

Select the check box under **Allow Any Query/Fragment** if you want to allow any query parameters or fragment in the resource.

 **Note:** When selected, no query parameter, fragment, or a combination of them is allowed in the path.

Select the check boxes under **TargetResource for SSO**, **TargetResource for SLO**, and **InError Resource** to indicate the resource types.

Click **Add**.


3. Repeat the previous step to add additional entries as needed.

Note that a more specific match is considered a better match and an exact match is considered the best match.

4. Click **Save**.

Manage partner redirect validation

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 via open redirections, PingFederate provides an option to validate `wreply` for SLO. Once enabled, the whitelist for the parameter is managed within the connection settings 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 in new installations by default.

For backward compatibility, PingFederate Upgrade Utility does not enable this option if it was not selected in the previous PingFederate installation. Although optional, it is strongly recommended to enable wreply validation for SLO and to specify the allowed domains and paths for each WS-Federation IdP connection to prevent unauthorized access (see [Specify a service URL \(WS-Federation\)](#) on page 358).

To reach this screen:

1. Click **Server Configuration** on the Main Menu.
2. Click **Redirect Validation** under System Settings.
3. Click **Partner Redirect Validation**.

To enable wreply validation for SLO:

1. Select the Enable wreply validation for SLO check box.
2. Click **Save**.

To disable wreply validation for SLO:

1. Clear the Enable wreply validation for SLO check box.
2. Click **Save**.

Security management

PingFederate provides built-in certificate management to handle SSL/TLS server security, as well as certificate signing and verification of SSO and other transactions, when required.

In addition, the server provides authentication capabilities for applications making use of secured system features, or for protocol features requiring management and validation of end-user password credentials.

This section covers:

- [Certificate management](#) on page 178
- [Authentication](#) on page 195



Note: This information is presented from the viewpoint of an administrative user with “Crypto Admin” permissions (see [Account management](#) on page 108).

Certificate management

Using the menu items under the **Server Configuration > Certificate Management** section, administrators may perform the following tasks:

- [Manage trusted certificate authorities](#) on page 178
- [Manage SSL server certificates](#) on page 179
- [Manage SSL client keys and certificates](#) on page 182
- [Manage digital signing certificates and decryption keys](#) on page 184
- [Manage certificate rotation settings](#) on page 189
- [Manage certificates from partners](#) on page 190
- [Configure certificate revocation](#) on page 191



Note: Access to **Certificate Management** is restricted to administrative users with the **Crypto Admin** administrative role (see [Account management](#) on page 108).

Manage trusted certificate authorities

You can import your federation partner's CA certificate or self-signed certificates into PingFederate's global trust list on the **Server Configuration > Trusted CAs** screen. If the Certificate Authority is not one of the major authorities, you may also need to import the certificate from the CA that signed the partner certificate.



Note: If a required CA certificate is already available in `cacerts` in the Java runtime, it is not necessary to import the same certificate into the PingFederate store.

To import a certificate

1. Click **Import**.
2. On the **Import Certificate** screen:
 - a) Click **Choose file** to locate and select the applicable certificate.
 - b) Select the storage facility of the certificate from the **Cryptographic Provider** list.
 - **HSM:** the integrated hardware security module (HSM).
 - **Local Trust Store:** the local trust store managed by PingFederate.

Applicable and visible only when PingFederate is integrated with an HSM from Thales in hybrid mode. (For more information, see [Supported hardware security modules](#) on page 51 .)
 - c) Click **Next**.
3. On the **Summary** screen, review the selected certificate and then click **Save** to complete the process.

To review a certificate

1. Select the applicable certificate by its serial number.
2. Review the selected certificate in the pop-up browser window.

When finished, close the pop-up window.
3. On the **Certificate Management** screen, click **Cancel** to go back to the **Server Configuration** screen.

To export a certificate

1. Click **Export** under **Action** for the certificate you want to export.
2. On the **Export Certificate** screen, click **Next**.
3. On the **Export & Summary** screen, click **Export** to save the certificate file on your system.

When finished, click **Done**.
4. Click **Cancel** to go back to the **Server Configuration** screen.

To delete a certificate

1. Click **Delete** under **Action** for the certificate you want to delete.

To undo the deletion, click **Undelete**.
2. Click **Save** to confirm the removal of the certificate.

Manage SSL server certificates

Use the **Server Configuration > SSL Server Certificates** screen to establish and maintain the certificates presented for access to the PingFederate administrative console and for incoming HTTPS connections at runtime.

To create a new certificate


Establish the HTTPS server certificate, which may be presented to incoming HTTPS connections.

1. Click **Create New**.
2. On the **Create Certificate** screen, enter the requested information.

For information about each field, refer to the following table:

Field	Description
Common Name	The common name (CN) identifying the certificate.
Organization	The organization (O) or company name creating the certificate.

Field	Description
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. <ul style="list-style-type: none"> • HSM: the integrated hardware security module (HSM). • Local Trust Store: the local trust store managed by PingFederate. <p>Visible only when PingFederate is integrated with an HSM in hybrid mode. (For more information, see Supported hardware security modules on page 51 .)</p>
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-SHA256, SHA384 and SHA512; and ECDSA-SHA256, SHA384 and SHA512.)

 **Note:** When using PingFederate with the Thales nShield Connect HSM, it is not possible to use an elliptic curve (EC) certificate as an SSL server certificate.

Select **RSA** and an RSA signing algorithm from the **Key Algorithm** list and the **Signature Algorithm** list, respectively.

When finished, click **Next**.


3. On the **Summary** screen, review the certificate information and then click **Save** to complete process.

To import a certificate and its private key

1. Click **Import**.

2. On the **Import Certificate** screen:

a) Click **Choose file** to locate and select the applicable certificate.

 **Note:** When using PingFederate with the Thales nShield Connect HSM, it is not possible to use an elliptic curve (EC) certificate as an SSL server certificate.

Select a certificate that uses the RSA key algorithm.

b) Enter the certificate password.

c) Select the storage facility of the certificate from the **Cryptographic Provider** list.

- **HSM**: the integrated hardware security module (HSM).
- **Local Trust Store**: the local trust store managed by PingFederate.

Applicable and visible only when PingFederate is integrated with an HSM in hybrid mode.

d) Click **Next**.

3. On the **Summary** screen, review the certificate information and then click **Save** to complete process.

To review a certificate

1. Select the applicable certificate by its serial number.

2. Review the selected certificate in the pop-up browser window.



Note: If a certificate has been revoked, PingFederate indicates this problem in the certificate information window.

When finished, close the pop-up window.

3. On the **Certificate Management** screen, click **Cancel** to go back to the **Server Configuration** screen.

To activate a certificate

1. Click **Activate for Runtime Server**, **Activate for Admin Console**, or both under **Action** for the certificate you want to activate.

These choices are enabled only if you have created or imported more than one certificate. Otherwise, a single certificate is used for both the administrative console and runtime operations.

To undo the activation, click **Cancel**.

2. Click **Save** to complete the process.

To create a certificate-authority signing request

1. Click **Certificate Signing** under **Action** for the desired 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 screen to initiate the process.



Tip: 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 CA signing.

2. On the **Certify Certificate** screen, select the **Generate Certificate Signing Request (CSR)** option and then click **Next**.
3. On the **Generate CSR** screen, click **Export** to save the CSR file on your system.

When finished, click **Done**.

4. On the **Certificate Management** screen, click **Cancel** to go back to the **Server Configuration** screen.

To import a certificate-authority response

1. Click **Certificate Signing** under **Action** for the relevant certificate.
2. On the **Certify Certificate** screen, select the **Import CSR Response** option and then click **Next**.
3. On the **Import CSR Response** screen, click **Choose file** to locate and select the CSR response file.

When finished, click **Next**.

4. On the **Summary** screen, review the information and then click **Save** to complete process.

To export a certificate

1. Click **Export** under **Action** for the certificate you want to export.
2. On the **Export Certificate** screen, select the type of export and then click **Next**.
 - **Certificate Only:** an export of the certificate without its private key. This is default choice.
 - **Certificate and Private Key:** an export of the certificate with its private key.



Caution: This export contains the private key of the certificate. You must also enter an encryption password.

If the certificate is stored on an HSM, the **Certificate and Private Key** option does not apply.

3. On the **Export & Summary** screen, click **Export** to save the certificate file on your system.

When finished, click **Done**.

4. Click **Cancel** to go back to the **Server Configuration** screen.

To delete a certificate

1. Click **Delete** under **Action** for the certificate you want to delete.

This option does not appear if the certificate is in use. To enable deletion, add (if needed) and activate a different certificate for the administrative console and/or the runtime server. If the usage is not clear, click **Check Usage**.

To undo the deletion, click **Undelete**.

2. Click **Save** to confirm the removal of the certificate.

Manage SSL client keys and certificates

On the **Server Configuration > SSL Client Keys & Certificates** screen, you create and manage your authentication private keys and the certificates your server presents as a client in an outbound SSL/TLS transaction.

To create a new certificate

1. Click **Create New**.
2. On the **Create Certificate** screen, enter the requested information.

For information about each field, refer to the following table:

Field	Description
Common Name	The common name (CN) identifying 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. <ul style="list-style-type: none"> • HSM: the integrated hardware security module (HSM). • Local Trust Store: the local trust store managed by PingFederate. <p>Visible only when PingFederate is integrated with an HSM in hybrid mode. (For more information, see Supported hardware security modules on page 51 .)</p>
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-SHA256, SHA384 and SHA512; and ECDSA-SHA256, SHA384 and SHA512.)

When finished, click **Next**.

3. On the **Summary** screen, review the certificate information and then click **Save** to complete process.

To import a certificate and its private key

1. Click **Import**.
2. On the **Import Certificate** screen:
 - a) Click **Choose file** to locate and select the applicable certificate.
 - b) Enter the certificate password.

c) Select the storage facility of the certificate from the **Cryptographic Provider** list.

- **HSM**: the integrated hardware security module (HSM).
- **Local Trust Store**: the local trust store managed by PingFederate.

Applicable and visible only when PingFederate is integrated with an HSM in hybrid mode.

d) Click **Next**.

3. On the **Summary** screen, review the certificate information and then click **Save** to complete process.

To review a certificate

1. Select the applicable certificate by its serial number.
2. Review the selected certificate in the pop-up browser window.



Note: If a certificate has been revoked, PingFederate indicates this problem in the certificate information window.

When finished, close the pop-up window.

3. On the **Certificate Management** screen, click **Cancel** to go back to the **Server Configuration** screen.

To activate a certificate

1. Click **Activate** under **Action** for the desired certificate.

These choices are enabled only if you have created or imported more than one certificate.

To undo the activation, click **Cancel**.

2. Click **Save** to complete the process.

To create a certificate-authority signing request

1. Click **Certificate Signing** under **Action** for the desired 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 screen to initiate the process.



Tip: 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 CA signing.

2. On the **Certify Certificate** screen, select the **Generate Certificate Signing Request (CSR)** option and then click **Next**.
3. On the **Generate CSR** screen, click **Export** to save the CSR file on your system.

When finished, click **Done**.

4. On the **Certificate Management** screen, click **Cancel** to go back to the **Server Configuration** screen.

To import a certificate authority response

1. Click **Certificate Signing** under **Action** for the relevant certificate.
2. On the **Certify Certificate** screen, select the **Import CSR Response** option and then click **Next**.
3. On the **Import CSR Response** screen, click **Choose file** to locate and select the CSR response file.

When finished, click **Next**.

4. On the **Summary** screen, review the information and then click **Save** to complete process.

To export a certificate

1. Click **Export** under **Action** for the certificate you want to export.
2. On the **Export Certificate** screen, select the type of export and then click **Next**.
 - **Certificate Only**: an export of the certificate without its private key. This is default choice.

- **Certificate and Private Key:** an export of the certificate with its private key.



Caution: This export contains the private key of the certificate. You must also enter an encryption password.

If the certificate is stored on an HSM, the **Certificate and Private Key** option does not apply.

3. On the **Export & Summary** screen, click **Export** to save the certificate file on your system.

When finished, click **Done**.

4. Click **Cancel** to go back to the **Server Configuration** screen.

To delete a certificate

1. Click **Delete** under **Action** for the certificate you want to delete.

This option does not appear if the certificate is in use. To enable deletion, add (if needed) and activate a different certificate for the administrative console and/or the runtime server. If the usage is not clear, click **Check Usage**.

To undo the deletion, click **Undelete**.

2. Click **Save** to confirm the removal of the certificate.

Manage digital signing certificates and decryption keys

You can use PingFederate to create and maintain your server's signing certificates, which you may use to sign outgoing requests, responses, assertions, and access tokens. The same type of certificate is also used for decryption.



Caution: For best security, we recommend using separate certificates for signing and decryption.

After creating your certificates, if they are left to be self-signed certificates, you can optionally enable automatic certificate rotation.

To create a new certificate

1. Click **Create New**.
2. On the **Create Certificate** screen, enter the requested information.

For information about each field, refer to the following table:

Field	Description
Common Name	The common name (CN) identifying 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. <ul style="list-style-type: none"> • HSM: the integrated hardware security module (HSM). • Local Trust Store: the local trust store managed by PingFederate. <p>Visible only when PingFederate is integrated with an HSM in hybrid mode. (For more information, see Supported hardware security modules on page 51 .)</p>
Key Algorithm	A cryptographic formula used to generate a key. PingFederate uses either of two algorithms, RSA or EC.

Field	Description
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-SHA256, SHA384 and SHA512; and ECDSA-SHA256, SHA384 and SHA512.)

When finished, click **Next**.

3. On the **Summary** screen, review the certificate information and then click **Save** to complete process.

To import a certificate and its private key


1. Click **Import**.
2. On the **Import Certificate** screen:
 - a) Click **Choose file** to locate and select the applicable certificate.
 - b) Enter the certificate password.
 - c) Select the storage facility of the certificate from the **Cryptographic Provider** list.
 - **HSM**: the integrated hardware security module (HSM).
 - **Local Trust Store**: the local trust store managed by PingFederate.

Applicable and visible only when PingFederate is integrated with an HSM in hybrid mode.

- d) Click **Next**.
3. On the **Summary** screen, review the certificate information and then click **Save** to complete process.

To review a certificate

1. Select the applicable certificate by its serial number.
2. Review the selected certificate in the pop-up browser window.

 **Note:** If a certificate has been revoked, PingFederate indicates this problem in the certificate information window.

When finished, close the pop-up window.

3. On the **Certificate Management** screen, click **Cancel** to go back to the **Server Configuration** screen.

To activate a certificate


1. Click **Activate** under **Action** for the desired certificate.
These choices are enabled only if you have created or imported more than one certificate.


To undo the activation, click **Cancel**.

2. Click **Save** to complete the process.

To create a certificate-authority signing request

1. Click **Certificate Signing** under **Action** for the desired 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 screen to initiate the process.

 **Tip:** 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 CA signing.

2. On the **Certify Certificate** screen, select the **Generate Certificate Signing Request (CSR)** option and then click **Next**.
3. On the **Generate CSR** screen, click **Export** to save the CSR file on your system.

When finished, click **Done**.

4. On the **Certificate Management** screen, click **Cancel** to go back to the **Server Configuration** screen.

To import a certificate authority response

1. Click **Certificate Signing** under **Action** for the relevant certificate.
2. On the **Certify Certificate** screen, select the **Import CSR Response** option and then click **Next**.
3. On the **Import CSR Response** screen, click **Choose file** to locate and select the CSR response file.

When finished, click **Next**.

4. On the **Summary** screen, review the information and then click **Save** to complete process.

To export a certificate

1. Click **Export** under **Action** for the certificate you want to export.
2. On the **Export Certificate** screen, select the type of export and then click **Next**.
 - **Certificate Only**: an export of the certificate without its private key. This is default choice.
 - **Certificate and Private Key**: an export of the certificate with its private key.



Caution: This export contains the private key of the certificate. You must also enter an encryption password.

If the certificate is stored on an HSM, the **Certificate and Private Key** option does not apply.

If the certificate was provided by your partner for signature verification or encryption, the **Certificate and Private Key** option also does not apply.

3. On the **Export & Summary** screen, click **Export** to save the certificate file on your system.

When finished, click **Done**.

4. Click **Cancel** to go back to the **Server Configuration** screen.

To delete a certificate

1. Click **Delete** under **Action** for the certificate you want to delete.

This option does not appear if the certificate is in use. To enable deletion, add (if needed) and activate a different certificate for the administrative console and/or the runtime server. If the usage is not clear, click **Check Usage**.

To undo the deletion, click **Undelete**.

2. Click **Save** to confirm the removal of the certificate.

Certificate rotation

PingFederate supports automatic certificate rotation for self-signed certificates created for the purpose of signing SAML requests, responses, and assertions, or XML decryption for Browser SSO and WS-Trust STS transactions on a per-certificate basis. This optional feature greatly reduces the cost of managing self-signed certificates.



Note: 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 twenty-five and ten percent of the original lifetime of the current certificate. 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 shutdown when the **Creation Buffer** threshold is reached for a given certificate, a new key pair and a new certificate is created if PingFederate is restarted during the rotation window.

Although optional, it is recommended that you turn on email notifications for certificate events in the **Server Configuration > Server Settings > Runtime Notification** screen. When configured, PingFederate notifies the configured recipient when a new certificate becomes available and when it is activated. Depending on the role of the certificate, you can update your partner accordingly.

Connection and federation metadata


Certification rotation is a per-certificate configuration. When certificate rotation is enabled for a certificate and a new certificate using a new key pairs becomes available, PingFederate deploys the new certificate to all active connections using that certificate. The actions taken by PingFederate vary depending on the role of the certificate.

Email notifications

Although optional, it is recommended that you turn on email notifications for certificate events in the **Server Configuration > Server Settings > Runtime Notification** screen. When configured, PingFederate notifies the configured recipient when a new certificate becomes 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 Browser SSO (SAML and WS-Federation) 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 becomes available, PingFederate performs the following tasks for all SAML 2.0 connections using the same decryption key:

- Push the current decryption key from primary to secondary.
- Place the new certificate as the primary decryption key.
- Update the decryption key with the new certificate in the metadata.
- Start 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 send email 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 supports the RSA key algorithm only. When EC (elliptic curve) is selected as the **Key Algorithm** value on the **Certificate Rotation** screen, 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 for you before the **Activation Buffer** threshold is reached.

Federation metadata for Browser SSO connections

PingFederate updates the metadata for the applicable Browser SSO connections as soon as a new certificate becomes available.

To ensure that your partners are aware of the new certificate, you can provide the partners their respective federation metadata URL or metadata export.

Metadata by URL

PingFederate runtime engine provides an endpoint (`/pf/federation_metadata.ping`) to return metadata for Browser SSO connections. An SP or an IdP is identified by its entity IDs using the **PartnerSpId** query parameter or the **PartnerIdpId** query parameter, respectively, as illustrated by the following examples.

Partner	Federation metadata URL to be given to the partner
An SP partner with an entity ID of SP1.	<code>https://www.example.com:9031/pf/federation_metadata.ping?PartnerSpId=SP1</code>
An IdP partner with an entity ID of IdP1.	<code>https://www.example.com:9031/pf/federation_metadata.ping?PartnerIdpId=IdP1</code>

Note that 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 the **Server Configuration > Cluster Management** screen when the new certificate becomes available, such that the federation metadata for these connections are 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 **Manage All** connections management screen or the **Server Configuration > Metadata Export** wizard.

- 📄 **Note:** PingFederate does not deploy new certificates or update metadata for inactive connections.

WS-Trust STS connections

For connections with only the WS-Trust 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 Browser SSO profile. Your partner can reuse the certificate from the metadata (by URL or manual export) and apply it to its STS configuration.

Managed SP connection to PingOne and signing certificate

PingFederate® automatically rotates the signing certificate used by the managed SP connection.


- 📄 **Note:** A managed SP connection to PingOne is a connection created by the initial setup wizard or the **Server Configuration > Connect to PingOne** wizard in PingFederate 8.0 (or higher).

The certificate rotation settings are as follow:

Field	Values
Creation Buffer (days)	90
Activation Buffer (days)	30
Validity (days)	1095
Key Algorithm	RSA
Key Size	2048
Signature Algorithm	SHA256withRSA

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 the **Server Configuration > Cluster Management** screen. When the process completes, PingFederate (the console node) uploads the new signing certificate to PingOne. If the console node is shutdown, it will create a new signing certificate provided that it is restarted during the rotation window.


Optionally, PingFederate can be configured to send reminders weeks ahead of expiry on the **Server Configuration > Server Settings > Runtime Notifications** screen.

 **Note:** If you prefer to manually rotate the signing certificate, disable automatic certificate rotation. Note that PingFederate still automatically uploads your new signing certificate to PingOne to maintain SSO continuity through the managed SP connection.

Manage certificate rotation settings

You manage certificate rotation settings for self-signed certificates in the **Server Configuration > Signing & Decryption Keys & Certificates** screen.

1. On the **Certificate Management** screen, click **Certificate Rotation** next to the applicable certificate.

 **Note:** Certificate rotation is only available to self-signed certificates.

2. On the **Enable Certificate Rotation** screen, select the check box to turn on certificate rotation for this certificate (*the current certificate*).

If you want to turn off certificate rotation for this certificate, clear the check box and click **Save**.

3. On the **Certificate Rotation** screen, modify the following fields as needed:

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 supports the RSA key algorithm only. When EC (elliptic curve) is selected as the **Key Algorithm** value on the **Certificate Rotation** screen, 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.

Signature Algorithm

The signing algorithm of the certificate. (RSA-SHA256, SHA384 and SHA512; and ECDSA-SHA256, SHA384 and SHA512.) The default value matches that of the current certificate.

4. On the **Certificate Rotation Summary** screen, review the rotation settings. Adjust as needed or click **Save** to turn on automatic certificate rotation for this certificate.

Manage certificates from partners

You receive certificates from partners for the following purposes:

- Signature verification: validating the digital signature in the incoming messages from your partners.
- Encryption: encrypting the outgoing messages before delivering them to your partners.
- Back-channel authentication: authenticating inbound (SOAP) messages from your partners by their client certificates.

You manage certificates from partner within a connection as follows:

- To manage certificates for the purpose of signature verification:
 - a) On the **Activation & Summary** screen, click **Signature Verification Certificate**.
 - b) On the **Signature Verification Certificate** screen, click **Manage Certificates**.
You can import, review, activate, export, or delete certificates on the **Manage Digital Verification Certificates** screen.
 - To manage certificates for the purpose of encryption:
 - a) On the **Activation & Summary** screen, click **Select XML Encryption Certificate**.
 - b) On the **Select XML Encryption Certificate** screen, click **Manage Certificates**.
You can import, review, activate, export, or delete certificates on the **Manage XML Encryption Certificates** screen.
 - To manage certificates for the purpose of inbound (SOAP) back-channel authentication:
 - a) On the **Activation & Summary** screen, click **SSL Verification Certificate**.
 - b) On the **SSL Verification Certificate** screen, click **Manage Certificates**.
You can import, review, activate, export, or delete certificates on the **Manage Verification Certificates** screen.
-  **Tip:** Depending on the use cases, your connection to the partner may not require signature verification, encryption, inbound (SOAP) back-channel authentication by client certificate, or any such combinations. If so, the **Activation & Summary** screen does not display the related administrative items.

To import a partner certificate

1. Click **Import**.
2. On the **Import Certificate** screen:
 - a) Click **Choose file** to locate and select the applicable certificate from your partner.
 - b) Select the storage facility of the certificate from the **Cryptographic Provider** list.
 - **HSM:** the integrated hardware security module (HSM).
 - **Local Trust Store:** the local trust store managed by PingFederate.
Applicable and visible only when PingFederate is integrated with an HSM from Thales in hybrid mode.
 - c) Click **Next**.
3. On the **Summary** screen, review the certificate information and then click **Save** to complete process.

To review a partner certificate

1. Select the applicable certificate by its serial number.
2. Review the selected certificate in the pop-up browser window.



Note: If a certificate has been revoked, PingFederate indicates this problem in the certificate information window.

When finished, close the pop-up window.

3. On the **Certificate Management** screen, click **Cancel** to go back to the **Server Configuration** screen.

To activate a partner certificate

1. Click **Activate** under **Action** for the desired certificate.

These choices are enabled only if you have created or imported more than one certificate.

To undo the activation, click **Cancel**.

2. Click **Save** to complete the process.

To export a partner certificate

1. Click **Export** under **Action** for the certificate you want to export.
2. On the **Export Certificate** screen, click **Next**.
3. On the **Export & Summary** screen, click **Export** to save the certificate file on your system.

When finished, click **Done**.

4. Click **Cancel** to go back to the **Server Configuration** screen.

To delete a partner certificate

1. Click **Delete** under **Action** for the certificate you want to delete.

This option does not appear if the certificate is in use. To enable deletion, add (if needed) and activate a different certificate for the administrative console and/or the runtime server. If the usage is not clear, click **Check Usage**.

To undo the deletion, click **Undelete**.

2. Click **Save** to confirm the removal of the certificate.

Configure certificate revocation

By default at runtime, PingFederate attempts to retrieve a CRL to verify that a signing certificate has not been revoked, whenever a CRL distribution-point URL is included within the certificate. Optionally, on the Manage Certificate Revocation screen you can enable and configure OCSP checking as the preferred verification method, depending on your requirements. (For more information, see [Certificate validation](#) on page 76.)

OCSP can be used in place of CRL checking, or CRLs can be retained as a backup method (for failover).




Note: When OCSP is enabled, CRL checking is not done independently—only as a failover option for one or more OCSP failure conditions.


On the **Server Configuration > Certificate Revocation Checking** screen, modify settings for CRL checking and OCSP, as needed.


Field Descriptions (For OSCP Checking)

Field/Selection	Description
Enable OCSP	Turns on OCSP certificate-revocation checking.
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, as needed).

Field/Selection	Description
Do NOT allow Responder to use cached responses	<p>When unchecked (the default), 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,” below).</p> <p>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.</p> <p> Important: Making this selection may slow down OCSP response time for a request and will increase general processing overhead at the Responder site.</p>
This Update Grace Period	For the response to be considered 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).
Next Update Grace Period	If the response includes a <nextUpdate> timestamp indicating when updated certificate statuses will be 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.
Responder Timeout	The allowable response time before the OCSP Responder URL is considered unavailable and processing continues (see “OCSP Responder is Unavailable,” below).
Certificate is Unknown	The certificate does not fall under the purview of the CA associated with the OCSP Responder. The drop-down choices indicate whether an unknown certificate is to be considered valid or not, or whether to try CRL checking.
OCSP Responder is Unavailable	Indicates what action to take if the Responder cannot be reached.
OCSP Responder Returns Error	Indicates what action to take if the Responder returns an error.
Proxy Settings	If OCSP messaging 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 CRL checking, when CRL is enabled for failover.

Field Descriptions (For CRL Checking)


Field/Selection	Description
Enable CRL Checking	<p>Enables CRL revocation checking (the default).</p> <p> 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.</p>
Treat Unretrievable CRLs as Revoked	<p>If checked, PingFederate immediately aborts the processing associated with the certificate.</p> <p>If unchecked, the server treats the certificate as valid but continues trying to retrieve the CRL.</p>
Next Retry on Resolution Failure	Specifies the number of minutes the server waits before trying to retrieve a CRL when the previous attempt failed—applies only when the selection above (Treat Unretrievable CRLs as Revoked) is unchecked.
Next Retry on Next Update Expiration	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.

Field/Selection	Description
	 Note: Certain actions in the administrative console, such as saving changes to an IdP adapter instance, reset the CRL cache. When it happens, PingFederate requests new CRLs for subsequent transactions as needed.
Verify CRL Signature	When checked (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 Manage trusted certificate authorities on page 178).
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.


Transition to an HSM

Starting with PingFederate® 8.3, administrators may enable the HSM hybrid mode, which provides the choice to store each relevant key and certificate on a hardware security module (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 to mirror the setup.


When all relevant keys and certificates are stored on the HSM, administrators may 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 if such keys and certificates exist on the local trust store.

1. Install and configure the HSM client and the existing PingFederate environment (see [Supported hardware security modules](#) on page 51)

 **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.

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 may 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 that are stored on your HSM and the local trust store.


 **Important:** When making changes to keys and certificates, you may need to coordinate with your partners. For more information, see [Digital signing policy coordination](#).

2. Create a new SSL server certificate on your HSM and activate it for the administrative console and the runtime server on the **Server Configuration > SSL Server Certificates** screen.

You may 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 179.

3. Create new digital signing certificates and decryption keys on the **Server Configuration > Signing & Decryption Keys & Certificates** screen and reconfigure connections or configuration items to use the new certificates and keys from your HSM.

 **Tip:** You may use **Check Usage** to locate the applicable connections or configuration items.


For configuration steps, see [Manage digital signing certificates and decryption keys](#) on page 184.

4. If your connections support outbound (SOAP) back-channel authentication by client certificates, create new SSL client certificates on the **Server Configuration > SSL Client Keys & Certificates** screen and reconfigure connections to use the new certificates from your HSM.

 **Tip:** You may use **Check Usage** to locate the applicable connections or configuration items.


For configuration steps, see [Manage SSL client keys and certificates](#) on page 182.

5. If you are transitioning to a Thales HSM, export the trusted CA certificates from the local trust store and import them to your HSM on the **Server Configuration > Trusted CAs** screen and reconfigure configuration items to use the new certificates and keys from your HSM.

 **Tip:** You may use **Check Usage** to locate the applicable configuration items.

For configuration steps, see [Manage trusted certificate authorities](#) on page 178.

6. If you are transitioning to a Thales HSM, for connections using the unanchored trust model, export the partner certificates 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](#).)

 **Tip:** You receive certificates from partners for the purposes of signature verification, encryption, or back-channel authentication. Partner certificates are managed in two buckets on a per-connection basis. While certificates for the purposes of signature verification and encryption are treated as one set of keys and certificates, client certificates for the purpose of inbound (SOAP) back-channel authentication are managed as another set of keys and certificates.


For configuration steps, see [Manage certificates from partners](#) on page 190.

Manage metadata URLs

SAML metadata URL streamlines the processes of establishing and maintaining SAML connections. If your partner provides SAML metadata by URL, you may use the metadata URL for the following scenarios:

- Create a new SAML connection using the metadata URL and associate the metadata URL with the new connection.
- Enable or disable automatic update from the associated metadata URL.
- Add or update the metadata URL associated with an existing SAML connection.
- Update an existing SAML connection using the metadata URL instantly.


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 subsequently for any of the aforementioned scenarios, PingFederate validates the digital signature using the previously stored certificate. If the signature is correct, the process carries on. If there is a digital signature error, PingFederate aborts the process and provides an error with a recommended course of action.

 **Tip:** These capabilities allow you to quickly create connections with InCommon participants, to update the connections automatically or manually as the InCommon participants update their metadata, and to do so securely knowing PingFederate only commits changes to your connections after the digital signatures of the signed metadata are validated. For more information about InCommon participants, please refer to www.incommon.org/participants.

Use the **Server Configuration > Metadata URLs** screen to add, update, review, or remove SAML metadata URLs provided by your partners.

Add a new metadata URL

1. In the **Metadata URLs** screen, click **Add New URL**.
2. In the **URL** screen, enter the name and the URL of the metadata.
3. In the **Certificate Summary** screen, review the certificate information.
 - a) If the metadata is not digitally signed (unsigned), click **Verify** to confirm that the unsigned metadata is reachable at the time of the configuration.
 - b) If the metadata is signed but the certificate is provided outside of the metadata, click **Import** to upload the verification certificate.
4. Review the configuration, and then click **Done** and **Save**.

-  **Tip:** When creating a new or updating an existing SAML connection, if you enter a new metadata URL, the metadata URL is automatically added to the system.

Update an existing metadata URL

1. In the **Metadata URLs** screen, click the name of the applicable metadata URL.
2. Access the **URL** screen to update the name, the URL, or both, and then click **Next**.
3. Access the **Certificate Summary** screen, and then 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.

If the metadata is signed but the certificate is provided outside of the metadata, click **Import** to upload the verification certificate.

4. Click **Next**.
5. Review the configuration, and then click **Done** and **Save**.

Review usage of or remove an existing metadata URL

1. In the **Metadata URLs** screen, click **Check Usage** for the applicable in-use metadata URL.

The pop-up window provides a list of connections associated with the selected metadata URL.

2. In the **Metadata URLs** screen, use the **Delete/Undelete** workflow to remove or cancel the removal request for the applicable metadata URL.



Note: You can only delete a metadata URL if it is not associated with any SAML connection.

3. Click **Save**.

Authentication

This portion of the Main Menu, under Server Configuration, allows administrators to manage Basic authentication to the PingFederate server, when needed, as well as authentication needs for secured protocol transactions.

Configure application authentication

When you use the SAML 2.0 Attribute Query profile as an SP, password security is required between the application requesting attributes and the SP PingFederate server. Basic authentication is also required for applications making calls to PingFederate's Connection Management Service and optional for the SSO Directory Service (see [Web Service interfaces](#) on page 477).

If you are using the SAML 2.0 Attribute Query profile as an SP, then the requesting application(s) at your site must authenticate to the PingFederate server (see [Attribute Query and XASP](#) on page 43 in the “Supported Standards” chapter of the *Get started with PingFederate* guide and [/sp/startAttributeQuery.ping](#) on page 454).

In addition, authentication is required to access PingFederate runtime data via JMX (see [Runtime monitoring using JMX](#) on page 148) or to make *SOAP* calls to the Connection Management Web Service. Authentication is optional for the SSO Directory Service (see [Web Service interfaces](#) on page 477).



Note: To help ensure network security, access to all of these services is deactivated when PingFederate is first installed.

On the **Server Configuration > Application Authentication** screen, administrators with the **Admin** administrative role can activate and configure authentication for the following services:

- Attribute Query
- JMX
- SSO Directory Service

To activate and configure authentication for the Connection Management Service, the administrators must be granted all three administrative roles: **Admin**, **Crypto**, and **User Admin**.

To enable access to a service:

1. Click **Activate** for the Service under Action.

- Where required, enter an Id, Shared Secret, and Confirm Shared Secret for the service.

You and the application developer must agree to these values.

This step is optional for the SSO Directory Service; the Service can be active without requiring authentication (the default setting).

- Repeat the steps above for other Services, as needed.
- Click **Save**.

To change an application ID or password:

- Replace the existing information in the necessary field(s) and click **Save**.

To block access to an active service:

- Click **Deactivate** for the Service under Action and then click **Save**.



Tip: Although not accessible when deactivated, the Connection Management Service and the SSO Directory Service are still deployed by default as part of PingFederate. If your organization does not plan to use one or neither of these services, you may wish to remove the corresponding WAR file or files from the <pf_install>/pingfederate/server/deploy and <pf_install>/pingfederate/server/deploy2 directories, respectively:

```
pf-ws.war
pf-mgmt-ws.war
```

Manage password credential validators

PingFederate® provides an authentication mechanism using plug-in password credential validators (PCVs). This feature provides centralized credential validation for other PingFederate configurations.

Multiple PCV instances are used by the HTML Form Adapter, the HTTP Basic Adapter, and the Username Token Processor. Multiple PCV instances can be applied to an IdP adapter or token processor. If an earlier PCV instance fails to authenticate the end user, control is returned to the IdP adapter or the token processor so that the next PCV instance can be tried.

Besides IdP adapters and token processors, PCV instances can also be mapped into the OAuth persistent grants in the **OAuth Settings > Resource Owner Credentials Mapping** screen for OAuth use cases using the password grant type.

PingFederate is distributed with the following plug-in PCVs:

LDAP Username Password Credential Validator

Validates credentials based on an LDAP look-up in an organization's user-data store.

PingOne Directory Password Credential Validator

Validates credentials stored in PingOne 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.

- 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**.



Note: Automatic multi-connection error checking occurs by default whenever you access this screen. The intent is to verify that configured connections have not been adversely affected by changes made here.

If you experience noticeable delays in accessing this screen, you can optionally disable automatic connection validation on the **Server Configuration > Server Settings > System Options** screen.

Choose a credential validator type

The first step in this configuration is choosing the type of the password credential validator (PCV). Available types are determined by plug-in JAR files loaded in the `<pf_install>/pingfederate/server/default/deploy` directory. Several validator plug-ins are bundled with PingFederate®. Other plug-ins may be added periodically, available from the Ping Identity [Downloads](#) website.

1. Enter a name and an ID for the instance on the **Type** screen.
2. Select the type of the PCV from the list.
3. 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 screens.

Configure a credential validator instance

The configuration of a password credential validator varies depending on the credential validators deployed on your server. For credential validators bundled with PingFederate®, refer to one of the following topics:

- [Configure the LDAP Username Password Credential Validator](#) on page 197
- [Configure the Simple Username Password Credential Validator](#) on page 200
- [Configure the RADIUS Username Password Credential Validator](#) on page 200
- [Configure the PingOne Directory Password Credential Validator](#) on page 202

Configure the LDAP Username Password Credential Validator

The LDAP Username Password Credential Validator (PCV) verifies credentials using an organization's LDAP data store.


When an authentication error occurs, PingFederate® automatically parses the messages returned by PingDirectory™, Microsoft Active Directory (AD), or Oracle Directory Server (ODS) and categorize them with the following error conditions:

- Account disabled
- Account expired
- Account locked
- Invalid credentials
- Not permitted to logon at this time
- Not permitted to logon at this workstation
- Password expired
- Password policy violated
- User must reset password
- User not found

As needed, and when validating against an LDAP directory server other than PingDirectory, AD, or ODS, administrators can define custom message categorization by mapping specific error messages (with wildcard support) to the desired error conditions on the **Instance Configuration** screen.

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 PCV. OAuth-client developers may create custom experiences based on the error responses, which contain the error messages.

The HTML Form Adapter also uses the relevant error conditions to determine the LDAP password-change scenarios and to present the relevant messages to the end users.

 **Tip:** These customizable messages are stored in the `<pf_install>/pingfederate/server/default/conf/language-packs/pingfederate-messages.properties` file.

As needed, they may be localized using the PingFederate localization framework for an international audience.

1. Go to the **Server Configuration > Password Credential Validators** screen.


- 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**.

2. On the **Type** screen:

- a) Enter a name and an ID for the password credential validator instance.
- b) Select the type from the list.
- c) 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 screens.

3. On the **Instance Configuration** screen, override authentication error messages as needed.

-  **Tip:** This option may be required for an LDAP directory server other than PingDirectory, AD, 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 may use wildcard asterisks (*) to match messages returned from your LDAP directory server; for example:

```
*expired*
```


- c) Select a relevant error condition from the **Error** list.
- d) Click **Update** under **Action**.
- e) Repeat these steps for additional overrides as needed.


Use the **Edit**, **Update**, and **Cancel** workflow to make or undo a change to an existing value. Use the **Delete** and **Undelete** workflow to remove an existing value or cancel the removal request.

Click **Move up** and **Move down** to change their display order. The ordering does not affect runtime processing.

4. Select the LDAP data store and enter information into the required fields.

For more information about each field, refer to the following table:

Field	Description
LDAP Datastore (Required)	The LDAP data store configured in PingFederate. If you have not yet configured the server to communicate with the LDAP 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 <i>must</i> secure the data store connection to your AD LDAP server using LDAPS; AD requires this level of security to allow password changes.
Search Base (Required)	The location in the LDAP directory server from which the search begins.
Search Filter	The LDAP query to locate a user record.

Field	Description
(Required)	<p>If your use case requires the flexibility of allowing users to identify themselves using different attributes, you may include these attributes in your query. For instance, the following LDAP filter allows users to sign on using either the sAMAccountName or mail attribute value through the HTML Form Adapter:</p> <pre data-bbox="570 342 1341 369">((sAMAccountName=\${username}) (mail=\${username}))</pre> <p> Important: To ensure that your 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:</p> <ul data-bbox="651 499 1430 856" style="list-style-type: none"> • You can extend the PCV contract and fulfill the subject identifier through the HTML Form Adapter (see Extend the contract for the credential validator on page 202). • You can add a data source in the SP connection and fulfill the subject identifier through a data store query (see Configure attribute sources and user lookup on page 296). • 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 (see Apply policy contracts to authentication policies on page 221). <p>Similarly, 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[®] 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 via an instance of the Composite Adapter or an authentication policy. For more information, see Input User Id Mapping in Configure the Composite Adapter on page 445 or Incoming User ID in Specify an incoming user ID on page 220, respectively.</p>
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 (the default) indicates a search of the base object and the entire subtree within the base object distinguished name.
Case-Sensitive Matching	The option to enable case-sensitive matching between the LDAP error messages returned from the LDAP directory server and the Match Expression values specified on this screen. This check box is selected by default.
Advanced fields for self-service password reset (SSPR) through the HTML Form Adapter	
Display Name Attribute	<p>The LDAP attribute used for personalizing messages to the user.</p> <p>This field is applicable for all password reset types (other than None).</p> <p>The default value is <code>displayName</code>.</p>
Mail Attribute	<p>The LDAP attribute containing the email address, to which PingFederate sends email notifications to the requesting users.</p> <p>This field is required when the password reset type is Email One-Time Link or Email One-Time Password in any invoking HTML Form Adapter instances.</p>

Field	Description
	The default value is <code>mail</code> .
SMS Attribute	<p>The LDAP attribute containing the phone number, to which PingFederate sends text message notifications to the requesting users.</p> <p>This field is required when the password reset type is Text Message in any invoking HTML Form Adapter instances.</p> <p>There is no default value.</p>
PingID Username Attribute	<p>The LDAP attribute containing the username to use for PingID based password reset.</p> <p>This field is required when the password reset type is PingID in any invoking HTML Form Adapter instances.</p> <p>There is no default value.</p>

Configure the Simple Username Password Credential Validator

The Simple username password credential validator verifies credentials maintained by PingFederate®.



Note: This validator is best used for testing purposes or for an organization with few accounts.

1. Go to the **Server Configuration > Password Credential Validators** screen.

- 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**.

2. On the **Type** screen:

- Enter a name and an ID for the password credential validator instance.
- Select the type from the list.
- 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 screens.

3. On the **Instance Configuration** screen, click **Add a new row to 'Users'**.

4. Enter the username and the password.

Use the **Edit**, **Update**, and **Cancel** workflow to make or undo a change to an existing value. Use the **Delete** and **Undelete** workflow to remove an existing value or cancel the removal request.

Click **Move up** and **Move down** to adjust the order in which you want PingFederate to attempt credential authentication. PingFederate moves sequentially through the list until credential validation succeeds. The credential validation process fails when no match is found.


5. To complete the configuration:

- Click **Done** on the **Summary** screen.
- Click **Save** on the **Manage Credential Validator Instances** screen.

Configure the RADIUS Username Password Credential Validator

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 one-time password in one field) and two-step (challenge-response) authentication. Two-step authentication is supported in the HTML Form Adapter.

 **Tip:** RADIUS server messages are used by the HTML Form Adapter to determine the two-step authentication scenarios and to present a login screen to the end users.


1. Go to the **Server Configuration > Password Credential Validators** screen.
 - 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**.
2. On the **Type** screen:
 - a) Enter a name and an ID for the password credential validator instance.
 - b) Select the type from the list.
 - c) 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 screens.

3. On the **Instance Configuration** screen, configure one or more RADIUS servers.
 - a) Click **Add a new row to 'RADIUS Servers'**.
 - b) Enter information into the required fields.

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, you can 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 Protocol	The protocol used to authenticate to the RADIUS server. The available choices are Password Authentication Protocol (PAP) and Challenge Handshake Authentication Protocol (CHAP). Select the protocol expected by your RADIUS server. The default value is PAP .
Shared Secret	The password shared between PingFederate and the RADIUS server used to encrypt passwords.

 **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** under **Action**.
- d) Repeat these steps for additional overrides as needed.

Use the **Edit**, **Update**, and **Cancel** workflow to make or undo a change to an existing value. Use the **Delete** and **Undelete** workflow to remove an existing value or cancel the removal request.

Click **Move up** and **Move down** 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.

4. 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 attribute identifying the NAS (Network Access Server) originating the request for access. The default value is <code>PingFederate</code> .
Timeout	The maximum number of milliseconds before a connection timeout to the RADIUS server.
Retry Count	The number of times to retry a failed connection before moving to the next host

Configure the PingOne Directory Password Credential Validator

The PingOne® Directory Username Password Credential Validator verifies credentials stored in the PingOne directory.



Note: The PingOne Directory Password Credential Validator requires a PingOne Enterprise account. For more information, see [Manage PingOne Directory Users](#).

1. Go to the **Server Configuration > Password Credential Validators** screen.

- 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**.

2. On the **Type** screen:

- Enter a name and an ID for the password credential validator instance.
- Select the type from the list.
- 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 screens.

3. On the **Instance Configuration** screen, enter your account information in the **Client ID** and **Client Secret**.

For more information about each field, refer to the following table. All fields are required.

Field	Description
Client ID and Client Secret	The REST API Client ID and its secret of your PingOne account, see View or Renew Directory API Credentials in the <i>PingOne Enterprise Administration Guide</i> .



Advanced Fields

PingOne URL	The PingOne directory API. The default value is <code>https://directory-api.pingone.com/api</code> .
Authenticate by Subject URL	The relative path for user authentication. The default value is <code>/directory/users/authenticate?by=subject</code> .
Reset Password URL	The relative path for password reset. The default value is <code>/directory/users/password-reset</code> .
SCIM User URL	The relative path for searching users requesting password reset. The default value is <code>/directory/user</code> .
Connection Pool Size	The maximum size of the connection pool to PingOne directory. The default value is 100.

Extend the contract for the credential validator

In some cases, you might want to extend contracts of the password credential validator 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 the LDAP server, the PingOne® directory, or the RADIUS server.

-  **Tip:** 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.
-  **Tip:** 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.

Vendor specific RADIUS attributes can be made available by extending the RADIUS attribute dictionary. Copy the vendor-specific attribute dictionaries into the `pingfederate/server/default/conf/radius` directory. The format of the dictionaries must use the *FreeRADIUS* dictionary syntax (freeradius.org/radiusd/man/dictionary.html). Then edit the existing `dictionary` file to include each of them.

- Optional: On the **Extended Contract** screen, enter an attribute name and click **Add**.

Use the **Edit**, **Update**, and **Cancel** workflow to make or undo a change to an existing attribute. Click **Delete** to remove an existing attribute.

Finish the credential validator configuration

- To complete and save the configuration, click **Done** on the **Summary** screen and then **Save** on the **Manage Credential Validator Instances** screen.


Configure Active Directory domains or Kerberos realms

From the **Manage AD Domains/Kerberos Realms** screen, provide PingFederate® with a centralized configuration to authenticate users via the following IdP adapters or token processors:

- **PingFederate integrated Kerberos Adapter** – Using the built-in Kerberos Adapter with a configured Active Directory (AD) Domain allows a PingFederate IdP server to perform SSO to SP applications based on Kerberos tickets.
- **PingFederate integrated Kerberos Token Processor** – The built-in Kerberos Token Processor accepts and validates Kerberos tokens via 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.

Follow these steps to configure an AD domain or Kerberos realm:


1. Configure the AD environment to integrate with PingFederate (see [Configure the Active Directory environment](#) on page 204).
2. Click **Add Domain/Realm** to create an AD domain.

-  **Important:** Do not configure subdomains if the parent domain in the same forest has already been configured (see [Multiple-domain support](#) on page 203).

Click the name to edit an existing domain. 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 [How Domain and Forest Trusts Work](http://technet.microsoft.com/en-us/library/cc773178(v=ws.10).aspx) ([technet.microsoft.com/en-us/library/cc773178\(v=ws.10\).aspx](http://technet.microsoft.com/en-us/library/cc773178(v=ws.10).aspx)) from Microsoft.


-  **Note:** If you are configuring only one domain, then you also need to configure only one Service Principal Name (see [Configure the Active Directory environment](#) on page 204).

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 a separate domain and then map these adapter or token processor instances to your SP connections that authenticate using the integrated Kerberos Adapter, the integrated Kerberos Token Processor, or the (separately available) IWA Adapter.

Configure the Active Directory environment

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 has already been configured (see [Multiple-domain support](#) on page 203).

 **Note:** You must have *Domain Administrator* permissions to make the required changes.

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 the password be set with no expiration.
2. Use the Windows utility `setspn` to register 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:


<pf-idp.domain.name>

The canonical name of the PingFederate server

(For more information on "canonical name", see <https://tools.ietf.org/html/rfc2181#section-10>.)

<pf-server-account-name>


The domain account you want to use for Kerberos authentication

 **Note:** "HTTP" must be capitalized and followed by a forward-slash (/).

3. Verify that the registration was successful by executing the following command:

```
setspn -l <pf-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 end-users already authenticated must re-authenticate (close the browser or log off and back on) before attempting SSO.

Add a domain

Use the **Manage Domain/Realm** screen to 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.

Enter the required information based on the following table:




Field	Description
Domain/Realm Name	The fully-qualified domain or realm name. For example: <code>companydomain.com</code>
Domain/Realm Username	The ID for the domain or realm account name.
Domain/Realm Password	The password for the domain or realm account.
Domain Controller/Key Distribution Center Host Names	Specify the host name or IP address of your domain controller or KDC (for example, <code>dc01-yvr</code>), and then click Add . Repeat this step to add multiple servers.

Field	Description
(optional)	<p>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 <i>unless</i> the Suppress DC / Domain Concatenation check box is selected.</p> <p>If unspecified, PingFederate uses a DNS lookup.</p>
Suppress DC / Domain Concatenation	<p>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 anymore.</p> <p>This check box is not selected by default.</p>
Test Domain/Realm Connectivity	<p>Tests access to the domain controller or KDC from the administrative-console server.</p> <p>When a connection to any of the configured controllers/KDCs is successful, the message <code>Test Successful</code> appears. Otherwise, the test returns error messages near the top of the screen.</p> <p> Tip: For help resolving connectivity issues, select the Debug Log Output check box on the Manage Domain/Realm Settings screen, run the test again, and review the debug messages in the PingFederate server log.</p> <p>Note that this test stops at the first successful result when multiple domain controllers or KDCs are specified; therefore, not all servers are necessarily verified. Depending on the network architecture, the engine nodes deployed in a cluster may also establish connections differently. As a result, the engine nodes and the console node may connect to different domain controllers or KDCs.</p>

Manage domain connectivity settings

Use the **Manage Domain/Realm Settings** screen to change default security and logging settings for all configured Active Directory domains and Kerberos realms.

Optional: Change the default transport protocol, the debug option, the timeout value, and the number of retry attempts, as needed. For more information of each field, refer to the following table:


Field	Description
Force TCP	<p>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.</p> <p> Note: If you choose this option, ensure that you restart PingFederate® after saving the configuration.</p>
Debug Log Output	<p>When selected, sends verbose messages to the PingFederate server log for all interactions with the domain controllers or the Key Distribution Centers (KDCs).</p>
AD Domain Controller/Key Distribution Center Timeout (secs)	<p>Specifies the amount of time (in seconds) PingFederate waits for a network response from a domain controller or KDC. The default is 3 seconds.</p> <p> Note: This value applies to each attempt PingFederate makes to contact the domain controller or KDC.</p> <p> Note: The new timeout takes effect only after PingFederate is restarted, after you save the configuration.</p>
AD Domain Controller/Key Distribution Center Retries	<p>Specifies the number of times PingFederate tries contacting the domain controller or KDC. The default is 3 times.</p>

Configure self-service password reset and account unlock

PingFederate® offers self-service password reset (SSPR) for end users to recover their account in the event of forgotten password. Integrated into the HTML Form Adapter and Password Credential Validator (PCV) framework, end users can now reset their password via one of four different mechanisms:

- One-time link via email
- One-time password via email
- One-time password via text message
- PingID®

The SSPR capability relies on the LDAP PCV and the HTML Form Adapter to query the required attributes for the chosen reset mechanism. PingFederate supports PingDirectory™, Microsoft Active Directory, and Oracle Directory Server out-of-the-box. Custom PCV implementations may also be developed to offer the SSPR features for users stored in non-LDAP data sources. For more information, refer to the `ResettablePasswordCredential` interface in Javadoc.


 **Tip:** The Javadoc for PingFederate is located the `<pf_install>/pingfederate/sdk/doc` directory.

PingFederate also provides the capability for end users to unlock their account without submitting a ticket to the IT department. When enabled with self-service password reset, if an account is locked, a user can initiate an account unlock request at the sign-on screen. Through the HTML Form Adapter, PingFederate prompts the user to prove ownership of the account using the password reset flow.

Unlike password reset, 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 Active Directory. If the underlying data store is connected to an Oracle Directory Server, users can only unlock their account by changing their current password through the password reset flow.

1. On the **Server Configuration > Server Settings > Runtime Notifications** screen, configure an email server that PingFederate can use to send email notifications for password reset (for password reset via one of the two email options) and to notify the status of the requests (for all password reset mechanisms).
2. On the **Server Configuration > Data Stores** screen, create a new LDAP data store.

You can also reuse an existing LDAP data store connection.

 **Note:** When connecting to an Active Directory (AD) LDAP server, you *must* secure the data store connection using LDAPS; AD requires this level of security to allow password changes.

 **Note:** When connecting to PingDirectory or Oracle Directory Server, configure proxied authorization for the service account on the directory server (see [Configure proxied authorization](#) on page 162).

(For more information, see [Configure an LDAP connection](#) on page 161.)

3. On the **Server Configuration > Password Credential Validators** screen, create a new LDAP Username PCV instance.

You can also reuse an existing LDAP Username PCV instance. If so, follow [this sub step](#) to configure the related advanced fields.

- a) Select a data store, enter a search base, define a search filter, select the scope of search, and enable or disable case-sensitive matching.
- b) Click **Show Advanced Fields** to update fields related to self-service password reset.

Configuration items vary depending on the desired password reset type. Refer to the following table for more information.

Field	Description
Display Name Attribute	The LDAP attribute used for personalizing messages to the user. This field is applicable for all password reset types (other than None). The default value is <code>displayName</code> .

Field	Description
Mail Attribute	<p>The LDAP attribute containing the email address, to which PingFederate sends email notifications to the requesting users.</p> <p>This field is required when the password reset type is Email One-Time Link or Email One-Time Password in any invoking HTML Form Adapter instances.</p> <p>The default value is <code>mail</code>.</p>
SMS Attribute	<p>The LDAP attribute containing the phone number, to which PingFederate sends text message notifications to the requesting users.</p> <p>This field is required when the password reset type is Text Message in any invoking HTML Form Adapter instances.</p> <p>There is no default value.</p>
PingID Username Attribute	<p>The LDAP attribute containing the username to use for PingID based password reset.</p> <p>This field is required when the password reset type is PingID in any invoking HTML Form Adapter instances.</p> <p>There is no default value.</p>

(For more information, see [Configure the LDAP Username Password Credential Validator](#) on page 197.)

4. On the **IdP Configuration > Adapters** screen, create a new HTML Form Adapter instance.

You can also reuse an existing HTML Form Adapter instance. If so, follow from [this sub step](#) to configure your adapter instance to enable self-service password reset and account unlock.

- a) Select the LDAP PCV instance defined in the previous step as the credential validator.
- b) Optional: Update any default values or options.
- c) Select the **Allow Password Changes** check box.
- d) Select the desired password reset type.

Property	Description
Password Reset Type	<p>Select one of the following methods for self-service password reset.</p> <ul style="list-style-type: none"> • Email One-Time Link: users receive an email with a URL to reset their password. • Email One-Time Password: users receive an email with a one-time password (OTP) to reset their password. • PingID: users are prompted to follow the PingID authentication flow to reset their password. <p>Ensure the PingID Username Attribute field in the selected LDAP PCV instance is configured; otherwise, users will not be able to reset their password.</p> <ul style="list-style-type: none"> • Text Message: users receive a text message notification with an OTP to reset their password. <p>Ensure the SMS Attribute field in the selected LDAP PCV instance is configured; otherwise, users will not receive text message notification for password reset.</p> <p>This option also requires SMS provider settings.</p> <ul style="list-style-type: none"> • None: users cannot reset password through this HTML Form Adapter instance. This is the default selection.


- e) Select the **Account Unlock** check box if you want to enable self-service account unlock as well.

- f) Click **Show Advanced Fields** to review or modify default values related to self-service password reset.

For self-service password reset using PingID, you must download the settings file from the PingOne admin portal and upload the file to the **PingID Properties** advanced field. In addition, if you have not enabled PingFederate integration in your PingID account, you must do so as well. For more information, see [HTML Form Adapter advanced fields](#) on page 435.

- g) If you have chosen **Text Message** as the password reset type, click **Manage SMS Provider Settings** to configure the SMS provider through which PingFederate can send text message notifications to the users.

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.

 **Tip:** For additional information about each field or Twilio, please refer to support.twilio.com.


Once saved, these SMS provider settings apply to all services using text message notifications.

(For more information, see [Configure the HTML Form Adapter](#) on page 431.)

5. Optional: Customize and localize the on-screen messages, email messages or text message (see [Customizable user-facing screens](#) on page 126, [Customizable email notifications](#) on page 133 and [Customizable text message](#) on page 138).

You have now successfully created a new instance or modified an existing instance of the HTML Form Adapter with the self-service password reset and account unlock capabilities.

When users sign on through this adapter instance, they have the option to reset their password or unlock their account using the **Trouble Signing On** link, as illustrated in this screen capture:



Configure account lockout protection

Account lockout protection prevents user accounts from becoming locked 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.

Account lockout protection is enabled in many areas of PingFederate[®]; for example, the HTML Form Adapter, the Username Token Processor, the OAuth Resource Owner Password Credentials grant type, and the administrative console native authentication scheme.

Settings are stored in the `com.pingidentity.common.security.AccountLockingService.xml` configuration file, located in the `<pf_install>/pingfederate/server/default/data/config-store` directory.



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 have the options to use unique values for different instances of the adapter or the token processor.

In a PingFederate clustered environment, this protection is shared across nodes, which helps in scenarios where PingFederate is deployed behind a load balancing infrastructure without sticky sessions. (For more information on how the underlying Account Locking Service works in a clustered environment, see [Account Locking Service](#) on page 549).

1. Open the `com.pingidentity.common.security.AccountLockingService.xml` file.

If you have a PingFederate clustered environment, open this file on the console node.

- To review or modify the maximum number of failed attempts before a user is locked out for a time period, look for the `MaxConsecutiveFailures` element value.

 **Note:** The per-instance setting in the HTML Form Adapter and the Username Token Processor overrides this property.

- To review or modify the amount of time (in minutes) that a user is locked out when the `MaxConsecutiveFailures` threshold is reached, look for the `LockoutPeriod` element value.

2. Save the change.

3. Restart PingFederate.

4. If you have a PingFederate clustered environment, click **Replicate Configuration** on the **Server Configuration > Cluster Management** screen.

Note that settings stored on the `com.pingidentity.common.security.AccountLockingService.xml` file are applicable to all services that make use of account lockout protection.


Authentication policies

Authentication policies, an optional configuration in PingFederate®, 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. Furthermore, administrators can reuse an authentication policy by ending it with an authentication policy contract and mapping such policy contract to multiple connections.

Selectors

Authentication selectors provide a plug-in capability for PingFederate® to evaluate various conditions related to the requests. PingFederate comes bundled with a set of authentication selectors. For 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 IdP connections based on the `PartnerIdpId` parameter values provided in the 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 the `<pf_install>/pingfederate/sdk/doc` directory.

Manage authentication selectors

You manage authentication selectors in the **Selectors > Manage Authentication Selector Instances** screen.

- 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**.

Choose a selector type

1. Enter a name and an ID for the authentication selector.

2. Select the desired type of authentication selector from the list.
3. 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 screens.

Configure an authentication selector

The configuration of an authentication selector varies depending on the authentication selectors deployed on your server. For authentication selectors bundled with PingFederate®, refer to one of the following topics:

- [Configure the CIDR Authentication Selector](#) on page 210
- [Configure the Cluster Node Authentication Selector](#) on page 211
- [Configure the Connection Set Authentication Selector](#) on page 211
- [Configure the HTTP Header Authentication Selector](#) on page 212
- [Configure the HTTP Request Parameter Authentication Selector](#) on page 213
- [Configure the OAuth Scope Authentication Selector](#) on page 214
- [Configure the Requested AuthN Context Authentication Selector](#) on page 215

Configure 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 SSO request. 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 in one or more authentication policies to route internal requests to an instance of the Kerberos Adapter.

1. Go to the **Selectors** screen from the **IdP Configuration** or **SP Configuration** menu.
 - To configure a new instance, 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** under **Action**.



Note: You can only remove a selector instance if it is not deployed in any authentication policy.

2. On the **Type** screen:
 - a) Enter a name and an ID for the authentication selector.
 - b) Select the type from the list.
 - c) 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 screens.

3. On the **Authentication Selector** screen:
 - a) Click **Add a new row to 'Networks'**, enter a network range, and then click **Update**.
 - b) Repeat the previous step as needed for additional ranges.

Use the **Edit**, **Update**, and **Cancel** workflow to make or undo a change to an entry. Use the **Delete** and **Undelete** workflow to remove an entry or cancel the removal request.


- c) Enter a **Result Attribute Name** value.

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.

Examples

An 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`.

 **Tip:** 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.

An IPv6 network range

Enter `2001:db8::/123` to cover 32 IPv6 addresses, ranging from `2001:db8::` through `2001:db8::1f`.

4. To complete the configuration:

- a) Click **Done** on the **Summary** screen.
- b) Click **Save** on the **Manage Authentication Selector Instances** screen.

Configure the Cluster Node Authentication Selector

The Cluster Node Authentication Selector enables PingFederate® to choose configured authentication sources or another selector based on the PingFederate cluster node that is servicing the request in one or more authentication policies. For example, this selector allows you to choose whether or not Integrated Windows Authentication is attempted based on the PingFederate cluster node with which a Key Distribution Center is associated.

1. Go to the **Selectors** screen from the **IdP Configuration** or **SP Configuration** menu.

- To configure a new instance, 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** under **Action**.

 **Note:** You can only remove a selector instance if it is not deployed in any authentication policy.

2. On the **Type** screen:

- a) Enter a name and an ID for the authentication selector.
- b) Select the type from the list.
- c) 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 screens.

3. On the **Authentication Selector** screen, click **Next** to proceed to the next screen.

4. On the **Selector Result Values** screen, list the index numbers for each of the cluster nodes.

Repeat this step to specify additional entries as needed.

Use the **Edit**, **Update**, and **Cancel** workflow to make or undo a change to an entry. Click **Delete** to remove an entry.

When an instance of the Cluster Node Authentication Selector is deployed in an authentication policy, each specified value forms its authentication policy path.

5. To complete the configuration:

- a) Click **Done** on the **Summary** screen.
- b) Click **Save** on the **Manage Authentication Selector Instances** screen.


Configure 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 SP connection used in an SSO request and SP connections

configured within PingFederate. This selector allows you to override connection authentication selection on an individual connection basis in one or more authentication policies.

1. Go to the **Selectors screen from the **IdP Configuration** or **SP Configuration** menu.**

- To configure a new instance, 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** under **Action**.

 **Note:** You can only remove a selector instance if it is not deployed in any authentication policy.

2. On the **Type screen:**

- a) Enter a name and an ID for the authentication selector.
- b) Select the type from the list.
- c) 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 screens.

3. On the **Authentication Selector screen:**

- a) Click **Add a new row to 'Connections'** under **Action**.
- b) Select an SP connection from the list and click **Update**.

At runtime, the selector compares the target SP connection in the SSO request to SP connections configured here. If a match is found, the selector returns a result value of `Yes`.

- c) Repeat the previous step as needed for additional connections.

There is no priority given to the order in which the connections appear.


Use the **Edit**, **Update**, and **Cancel** workflow to make or undo a change to an entry. Use the **Delete** and **Undelete** workflow to remove an entry or cancel the removal request.

4. To complete the configuration:

- a) Click **Done** on the **Summary** screen.
- b) Click **Save** on the **Manage Authentication Selector Instances** screen.


Configure 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. 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:** We do not recommend using this selector to determine whether, or not, an authentication source with a higher level of assurance should be bypassed because HTTP request headers could potentially be forged.

1. Go to the **Selectors screen from the **IdP Configuration** or **SP Configuration** menu.**

- To configure a new instance, 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** under **Action**.

 **Note:** You can only remove a selector instance if it is not deployed in any authentication policy.

2. On the **Type screen:**

- a) Enter a name and an ID for the authentication selector.
- b) Select the type from the list.
- c) 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 screens.

3. On the **Authentication Selector screen:**


- a) Click **Add a new row to 'Results'** under **Action**.
- b) Enter an expression for use when inspecting the HTTP header value of the target HTTP header under **Match Expression**, and then click **Update**.

Repeat as needed to define additional expressions.

Use the **Edit**, **Update**, and **Cancel** workflow to make or undo a change to an entry. Use the **Delete** and **Undelete** workflow to remove an entry or cancel the removal request.

Wildcard entries are allowed. For example, the following entries identify the most common browsers:

Browser	Expression
Chrome	*Chrome*
Firefox	*Firefox*
iPhone	*iPhone*
iPad	*iPad*
Internet Explorer	*MSIE*
Safari	*Safari*

 **Tip:** For information about Internet Explorer 11 (or higher), see [User-agent string changes](https://msdn.microsoft.com/library/hh869301.aspx) from Microsoft (msdn.microsoft.com/library/hh869301.aspx).

- c) Enter the type of HTTP header you want the selector to inspect in the **Header Name** field; for example, `User-Agent`.

This field is not case-sensitive.

- d) Optional: Clear the **Case-Sensitive Matching** check box to disable case-sensitive matching between the HTTP header values from the requests and the **Match Expression** values specified on this screen.

The **Case-Sensitive Matching** check box is selected by default.


At runtime, the selector compares the HTTP header to these wildcard expressions. If a match is found, the selector returns a result value of `Yes`.

4. To complete the configuration:

- a) Click **Done** on the **Summary** screen.
- b) Click **Save** on the **Manage Authentication Selector Instances** screen.

Configure 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. 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 SSO request.

 **Important:** We do not recommend using this selector to determine whether, or not, an authentication source with a higher level of assurance should be bypassed because query parameters could potentially be forged.

1. Go to the **Selectors screen from the **IdP Configuration** or **SP Configuration** menu.**

- To configure a new instance, 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** under **Action**.

 **Note:** You can only remove a selector instance if it is not deployed in any authentication policy.

2. On the **Type** screen:

- Enter a name and an ID for the authentication selector.
- Select the type from the list.
- 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 screens.

3. On the **Authentication Selector** screen:

- Enter the request parameter name in the **HTTP Request Parameter Name** field.

The HTTP request parameter name is case-sensitive.

- Optional: Clear the **Case-Sensitive Matching** check box 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** screen.


The **Case-Sensitive Matching** check box is selected by default.

4. On the **Selector Result Values** screen, enter a request parameter value.

Wildcard entries are allowed. Note that a more specific match is considered a better match and an exact match is considered the best match.


Repeat as needed to specify additional entries.

Use the **Edit**, **Update**, and **Cancel** workflow to make or undo a change to an entry. Click **Delete** to remove an entry.



Result Values	Action
alpha	Edit Delete
beta	Edit Delete
staging	Edit Delete

When an instance of the HTTP Request Parameter Authentication Selector is deployed in an authentication policy, each specified value forms its own authentication policy path.



Action	Result	Action
Selector - (Selector)	alpha	-SELECT-
	beta	-SELECT-
	staging	-SELECT-

5. To complete the configuration:

- Click **Done** on the **Summary** screen.
- Click **Save** on the **Manage Authentication Selector Instances** screen.

Configure 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 (OAuth AS).

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 rather than username and password.

1. Configure one or more scopes in **OAuth Settings > Authorization Server Settings** screen if you have not already done so.
2. Go to the **Selectors** screen from the **IdP Configuration** or **SP Configuration** menu.

- To configure a new instance, 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** under **Action**.


 **Note:** You can only remove a selector instance if it is not deployed in any authentication policy.

3. On the **Type** screen:
 - a) Enter a name and an ID for the authentication selector.
 - b) Select the type from the list.
 - c) 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 screens.

4. On the **Authentication Selector** screen, select the required scopes.

At runtime, the OAuth Scope Authentication Selector compares the requested OAuth scopes to the scopes selected here. All of the selected scopes must match for the selector to return a result value of **Yes**.

 **Important:** This selector matches only scopes from OAuth authorization requests to the authorization endpoint (`/as/authorization.oauth2`). SAML SSO requests do not match this authentication selector's criteria and result in a returned result value of **No**. Therefore, if you are using this selector and selectors specific to SAML connections, place this selector first in the mapping list so that it takes precedence for OAuth without disrupting selector logic on SAML connections.

5. To complete the configuration:
 - a) Click **Done** on the **Summary** screen.
 - b) Click **Save** on the **Manage Authentication Selector Instances** screen.


Configure the Requested AuthN Context Authentication Selector

The Requested AuthN Context Authentication Selector enables PingFederate® to choose configured authentication sources or other selectors based on the authentication context (or contexts) requested by an SP for Browser SSO requests or an RP for OAuth with OpenID Connect use cases in one or more authentication policies.

For Browser SSO, this authentication selector works in conjunction with SP connections via 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).

1. Go to the **Selectors** screen from the **IdP Configuration** or **SP Configuration** menu.

- To configure a new instance, 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** under **Action**.

 **Note:** You can only remove a selector instance if it is not deployed in any authentication policy.

2. On the **Type** screen:
 - a) Enter a name and an ID for the authentication selector.
 - b) Select the type from the list.

- c) 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 screens.

3. On the **Authentication Selector** screen, select the **Add or Update AuthN Context Attribute** check box if you want to update the AuthN Context attribute value with the value specified in the **Selector Result Values** screen accordingly.

When selected (the default), the check box on this screen provides a means of either:

- Adding the value of the authentication context determined by the selector into the SAML assertion
- When applicable, replacing any value returned from the associated adapter instance with the selector-result value

4. On the **Selector Result Values** screen, specify the authentication contexts to be used as the criteria.

- a) Enter the exact (case-sensitive) parameter value under **Result Values** and click **Add**.

The value may include URIs defined in [Authentication Context for the OASIS Security Assertion Markup Language \(SAML\) 2.0](https://docs.oasis-open.org/security/saml/v2.0/saml-authn-context-2.0-os.pdf) (docs.oasis-open.org/security/saml/v2.0/saml-authn-context-2.0-os.pdf) or any other value agreed upon with the partner.



Note: The selector returns true when the specified value is an exact (case-sensitive) match to the authentication context (or one of the authentication contexts) requested by the SP or the RP.

- b) Add more values to differentiate criteria for authentication selection.

Use the **Edit**, **Update**, and **Cancel** workflow to make or undo a change to an entry. Click **Delete** to remove an entry.

When an instance of the Requested AuthN Context Authentication Selector is deployed in an authentication policy, each specified authentication context forms its authentication policy path.

5. To complete the configuration:

- a) Click **Done** on the **Summary** screen.
- b) Click **Save** on the **Manage Authentication Selector Instances** screen.

Policies

An authentication policy is a tree of authentication sources, selector instances, and authentication policy contracts, on which the decisions to route the request to a series of approved authentication sources with an optional authentication policy contract at the end or deny the request are based. Administrators can create one or more authentication policies to fulfill their authentication requirements.

IdP authentication policies apply to IdP Browser SSO requests, adapter-to-adapter requests, and browser-based OAuth authorization-code and implicit flows. SP authentication policies on the other hand apply to SP-initiated Browser SSO requests received by the PingFederate® SP server at the `/sp/startSSO.ping` endpoint.

The order of authentication policies matters because the policy engine in PingFederate evaluates policies from top to bottom.

At runtime, the policy engine derives an authentication tree from the applicable policies and either approves or denies a request.

Policy, paths, and authentication policy contracts

An authentication policy starts with either a selector or an authentication source. Authentication sources and most selectors have two results (**Success** or **Fail**, **Yes** or **No**). Each result forms a path.

A 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 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.

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 policy contract for different closed-ended paths (in one or more authentication policies) and fulfill them differently to suit the requirements. For IdP Browser SSO use cases, the same policy contract can also be mapped to multiple SP connections, thus enforcing the same set of authentication policies to those connections.

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 may also create multiple policies to suit their needs.

When a request arrives at PingFederate, the policy engine skips any closed-ended paths that are not applicable to the request, which may skip one or more policies if necessary. A path is considered not applicable to a request if the authentication policy contract at the end of the path (or the last authentication source in such path when it does not end with an authentication policy contract) is not mapped to the corresponding use case or is blocked by the virtual server ID included in the request. Note that virtual server IDs are not applicable to adapter-to-adapter mappings or OAuth uses cases.


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**). Depending on the policies, the policy engine may 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 no default authentication source can be found 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 both the `IdpAdapterId` query parameter and the `pfidpaid` cookie are presented, 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 let 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 that 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.

Define authentication policies

You manage authentication policies and settings in the **Policies > Manage Authentication Policies** screen.

1. Select the **Enable IdP Authentication Policies** check box to turn on authentication policies, which apply to IdP Browser SSO requests, adapter-to-adapter requests, and browser-based OAuth authorization-code and implicit flows.

This check box is only visible when the IdP role is selected in the **Server Configuration > Server Settings > Roles & Protocols** screen.

If you are still configuring policies and connections, clear this check box for now. When you are ready, select the option to activate the policies.

2. Select the **Enable SP Authentication Policies** check box to turn on authentication policies, which apply to SP-initiated Browser SSO requests received by the PingFederate® SP server at the `/sp/startSSO.ping` endpoint.



Note: SP authentication policies are *not* applicable to IdP Browser SSO requests, browser-based OAuth authorization-code and implicit flows, or adapter-to-adapter requests that PingFederate receives at the `/pf/adapter2adapter.ping` endpoint.

This check box is only visible when the SP role is selected in the **Server Configuration > Server Settings > Roles & Protocols** screen.

If you are still configuring policies and connections, clear this check box for now. When you are ready, select the option to activate the policies.

3. 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.
4. Create an authentication policy.



Tip: Think of authentication sources and selectors as checkpoints when implementing your authentication requirements.

- a) Select an authentication source (an IdP adapter or an IdP connection) or a selector from the **Action** list.

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-screen instructions to manage your rules.

For the PingID Adapter, IdP adapters developed using the `IdpAuthenticationAdapterV2` interface from the PingFederate SDK (including the bundled HTML Form Adapter), and SAML 2.0 IdP connections supporting the SP-initiated Browser SSO profile, you may specify a user ID to be passed in from an earlier-factor adapter. Click **Options** and follow the on-screen instructions to select the source and the attribute to be used as the incoming user ID.

All outcomes (including those based on rules) are displayed under **Result**. Each result forms a path.

- b) For each path, select an action from the list.

- If additional processing is required, repeat *step 4a*.
- If the path is extended from an authentication source and it is the end of the path, select **--Done--**, which marks this path a closed-ended path.

A 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.

If you need to reuse an authentication policy in multiple connections, select an authentication policy contract as the last action of a path, configure its contract fulfillment, and map the authentication policy contract to the applicable connections. Click **Contract Mapping** and follow the on-screen instructions to configure the fulfillment of the authentication policy contract.

If the authentication policy contract you need is not available, create a new authentication policy contract or edit an existing authentication policy contract in the **Policies > Authentication Policy Contracts** screen.

- If the path is extended from a selector and it is the end of the path without any prior authentication source, select **--Continue--**, which leaves this path as an open-ended path.

A 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.

5. Optional: Repeat [step 4](#) to create additional authentication policies.



Important: Order matters because the policy engine starts from the first policy and works its way down.

6. Optional: 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.

7. Click **Save**.

If you are not ready to deploy all the policies yet, clear the **Enable IdP Authentication Policies** check box, the **Enable SP Authentication Policies**, or both, before saving your policies.

Configure rules in authentication policies

An authentication source in an authentication policy has two results, **Fail** or **Success**, for which one of the following actions can be set:

- 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 (also terminating the authentication policy, making it a closed-ended path)

A 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.

PingFederate® supports more granular control through the use of rules in authentication policies. 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 such user is a contractor. Your organization mandates that all contractors must authenticate successfully against the OpenToken IdP adapter, followed by another IdP adapter (for example, 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. Locate the authentication source that you want to define additional successful results for further processing.
2. Click **Rules** underneath **Success** (result).

The PingFederate administrative console should display a **Rules** dialog.

3. Select an attribute from the previous authentication source in the **Attribute Name** list.
4. Select how PingFederate should compare the value that you are going to specify in the next step against the attribute value from the authentication source in the **Condition** 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 multi-value conditions when you want PingFederate to verify that 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. Enter the desired value to be compared against the attribute value from the authentication source in the **Value** field.
6. Enter a unique label in the **Result** field.
7. Optional: Click **Add** and repeat steps 3 to 6 to add another rule.
8. Optional: Click **Delete** when a specific rule is no longer required.
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.

By default, this check box is selected. When cleared, the policy engine treats the attempt as a failure when no rules return a match.

10. Click **Done** to close the **Rules** dialog.

The PingFederate administrative console brings you back to the **Manage Authentication Policies** screen. The authentication policy is updated with new successful results based on your 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**
- **Success - Contractors** (a new result based on the first rule)
- **Success - Senior executives** (a new result based on the second rule)
- **Success** (available only when the **Default to Success** check box is selected)

Select an action for each result.

Specify an incoming user ID

Some authentication sources, such as the PingID[®] Adapter, require a user ID to be passed in from an earlier-authentication step to perform multifactor authentication. Additionally, all IdP adapters developed using the `IdpAuthenticationAdapterV2` interface from the PingFederate SDK (including the bundled HTML Form Adapter) and SAML 2.0 IdP connections supporting the SP-initiated Browser SSO profile may support an incoming user ID.

For these use cases, use the **Options > Incoming User ID** dialog to specify the incoming user ID in your authentication policies.

In a nutshell, for the authentication source that needs an incoming user ID, you select the source and the attribute in the policy; the source can be any IdP adapter instance, an IdP connection that is capable of handling SP-initiated Browser SSO requests, or the original SAML 2.0 AuthnRequest. At runtime, PingFederate maps the attribute value as the user ID.

1. Locate the authentication source that you need to provide an incoming user ID, and then click **Options** (underneath the entry).

The **Incoming User ID** dialog should open.

2. Select the source of the attribute from the list.

You can choose from any IdP adapter instances and IdP connections that were placed ahead of the current authentication source in the same path or `Context`, which represents the original SAML 2.0 AuthnRequest at runtime.

3. Select an attribute from the list.

If you have selected `Context` from the **Source** list, select the `Requested User` attribute, which represents the optional `subject` identifier in the SAML 2.0 AuthnRequest at runtime.

If a request does not originate with an SAML 2.0 AuthnRequest (for example, an IdP-initiated SSO request or a WS-Federation SSO request), or if a SAML 2.0 AuthnRequest does not include the optional `subject` identifier, the policy engine does not pre-populate the user ID.

4. Click **Done** to close the **Incoming User ID** dialog.

The PingFederate administrative console brings you back to the **Manage Authentication Policies** screen.

Apply policy contracts to authentication policies

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 policy contract for different closed-ended paths (in one or more authentication policies) and fulfill them differently to suit the requirements. For IdP Browser SSO use cases, the same policy contract can also be mapped to multiple SP connections, thus enforcing the same set of authentication policies to those connections.

To apply an authentication policy contract to a policy, select an authentication policy contract as the last action of one or more paths and configure fulfillment for each contract.

1. Locate all closed-ended paths in the policy.

A 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.

A 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.

You cannot map an authentication policy contract to an open-ended path because the users have not yet been challenged by any authentication source.

Consider the following sample policy:

Action	Result	Action	Result	Action	Result	Action	Result
Test (Selector)	No	HTML Form (Adapter)	Fail				
Test (Selector)	No	HTML Form (Adapter)	Success	Retail (Selector)	Yes	PingID (Adapter)	Fail
Test (Selector)	No	HTML Form (Adapter)	Success	Retail (Selector)	Yes	PingID (Adapter)	Success
Test (Selector)	Yes						
Test (Selector)	No	HTML Form (Adapter)	Success	Retail (Selector)	Yes	PingID (Adapter)	Success

This policy has five paths:

- **Test (Selector) > No > HTML Form (Adapter) > Fail**
- **Test (Selector) > No > HTML Form (Adapter) > Success > Retail (Selector) > No**
- **Test (Selector) > No > HTML Form (Adapter) > Success > Retail (Selector) > Yes > PingID (Adapter) > Fail**
- **Test (Selector) > No > HTML Form (Adapter) > Success > Retail (Selector) > Yes > PingID (Adapter) > Success**
- **Test (Selector) > Yes**

The first four paths are closed-ended while the last path is open-ended.

- For each applicable closed-ended path, select the applicable authentication policy contract as the last action from the **Action** list.

It makes sense to select an authentication policy contract for the **PingID (Adapter) > Success** result, because the users have successfully met all your authentication requirements. Depending on your use case, you may also select an authentication policy contract for the **PingID (Adapter) > 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 workflow in the applicable connections later.

If the authentication policy contract you need is not available, create a new authentication policy contract or edit an existing authentication policy contract in the **Policies > Authentication Policy Contracts** screen.

- For each selected authentication policy contract, click **Contract Mapping** and configure the authentication policy contract using the **Manage Authentication Policies > Authentication Policy Contract Mapping** wizard.
 - Optional: Click **Add Attribute Source** in the **Attribute Sources & User Lookup** screen to configure data store queries.
 - In the **Contract Fulfillment** screen, fulfill the selected contract.

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 sample policy, 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.
 - Optional: In the **Issuance Criteria** screen, configure conditions to be validated before issuing an authentication policy contract (see [Define issuance criteria for authentication policy contract mapping](#) on page 222).
 - In the **Summary** screen, review your configuration, modify as needed, and then click **Done**.

The PingFederate® administrative console brings you back to the **Manage Authentication Policies** screen.

Define issuance criteria for authentication policy contract mapping

Use the **Issuance Criteria** screen to define criteria that PingFederate® can evaluate to determine whether to approve or reject requests for protected resource.



The configuration begins by selecting a source containing the attribute to be verified. Some sources, such as **Context** and **Mapped Attributes**, are common to almost all use cases. Other sources, such as **Assertion**, have dependency on the type of configuration, for which the administrative console automatically hides the irrelevant sources automatically. After selecting a source, choose an attribute to be verified. Depending on the selected source, attributes vary. For example, if **Account Linking** is the selected source, **Local User ID** is the available attribute. Finally, select the comparison method, specify the desired value, and enter an optional error message for scenarios where the condition fails.

When multiple criteria are configured through the user interface, all criteria must be met at runtime (evaluated as *true*) for the request to succeed; otherwise PingFederate rejects the request and returns an error message. If only a subset or a layer of the criteria is required, use one or more OGNL expressions to define them.

- Select a source containing the attribute to be verified.

For more information about the **Source** list, refer to the following table:



Source	Description
Adapter	Select to access attributes from any IdP adapter instance leading up to this authentication policy contract.
IdP Connection	Select to access attributes from any IdP connection leading up to this authentication policy contract.
JDBC, LDAP, or Custom (if configured)	Select to access attributes returned from a data source.
Mapped Attributes	Select to access the mapped attributes.


2. Select an attribute from the **Attribute Name** list.
3. Select a comparison method from the **Condition** list.
 -  **Note:** Use multi-value conditions when you want PingFederate to verify that an attribute contains or does not contain the specified value in its attribute list. Using a single-value condition when an attribute has multiple values causes the condition to evaluate consistently to false.
4. Enter an exact value in the text field.
 -  **Note:** If your use case requires more complex evaluations, such as partial matching, use one or more OGNL expressions to define them.
5. Optional: Enter a custom error message in the **Error Result** field that PingFederate returns when the criterion fails.

If you leave this field blank, PingFederate returns `ACCESS_DENIED` at runtime when the criterion fails.
6. Click **Add**.

Use the **Edit**, **Update**, **Cancel**, **Delete**, and **Undelete** links to manage the criterion.
7. Optional: Repeat these steps to specify additional criteria.

Define issuance criteria using OGNL expressions

1. Click **Show Advanced Criteria**.
 -  **Note:** Expressions must be enabled for the **Show Advanced Criteria** button to appear.
 -  **Tip:** If you have not already done so, you may enable OGNL expression by editing the `org.sourceid.common.ExpressionManager.xml` file in the `<pf_install>/pingfederate/server/default/data/config-store` directory. 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 on the **Server Configuration > Cluster Management** screen, and restart all nodes.
2. Specify a criterion.
 - a) Enter an OGNL expression into the text field.
 - b) Enter a custom error message in the **Error Result** field.
 -  **Note:** If the OGNL expression resolves to a string value (rather than *true* or *false*), the returned value overrides the custom error message.
 - c) Click **Add**.
 - d) Click the **Test** link to input values and test the resulting output for the expression.

Use the **Edit**, **Update**, **Cancel**, **Delete**, and **Undelete** links to manage the OGNL expression.
3. Optional: Repeat these steps to specify another OGNL expression.

If multiple OGNL expressions are specified, they all must pass in order for the request to succeed. Alternatively, you can combine multiple criteria in one OGNL expression.

Map a policy contract to multiple connections

The last step to reuse an authentication policy in multiple SP connections is to map the authentication policy contract into the applicable SP connections.

Generally speaking, for IdP Browser SSO use cases, if you have selected authentication policy contracts in your authentication policies, you must map the authentication policy contracts to the applicable SP connections.

1. Select the applicable SP connection from the **IdP Configuration** menu.
2. In the **Activation & Summary** screen, click **Authentication Source Mapping**.

3. Click **Map New Authentication Policy** and follow the on-screen instructions to map the authentication policy contract into the SP connection.

Configure an SP authentication policy for users from one IdP

For SP-initiated Browser SSO requests received at the `/sp/startSSO.ping` application endpoint, PingFederate® provides the optional capability to enforce additional authentication requirements through the use of SP authentication policies.

When you enable SP authentication policies to enforce additional authentication requirements, you are configuring policies for the policy engine in the PingFederate SP server to find an applicable SP adapter instance to access your target applications. Therefore, you must configure the target applications to provide the `SpSessionAuthnAdapterId` parameter, the `TargetResource` parameter, or both, when submitting their SP-initiated SSO requests to the `/sp/startSSO.ping` application endpoint. If you prefer to provide the `TargetResource` parameter *without* the `SpSessionAuthnAdapterId` parameter, you must configure one or more entries in the **SP Configuration > Target URL Mapping** screen to map the `TargetResource` values to the applicable SP adapter instances.



Note: When SP authentication policies are turned on, they are applicable only to SP-initiated Browser SSO requests received at the `/sp/startSSO.ping` application endpoint and not unsolicited SSO requests (namely IdP-initiated Browser SSO requests via SAML 1.x or 2.0).

It is therefore worth noting that SP authentication policies are *not* applicable to IdP Browser SSO requests, browser-based OAuth authorization-code and implicit flows, or adapter-to-adapter requests that PingFederate receives at the `/pf/adapter2adapter.ping` endpoint.

Consider the following example:

You are tasked to create an IdP connection to Alpha, which passes two attributes in its assertions, `SAML_SUBJECT` and `samlEmail`, on your PingFederate SP server. You are also asked to enforce multifactor authentication 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 SSO profile.


You have already created an SP adapter instance using the **SP Configuration > Adapters** wizard and completed the last-mile integration with the target application. The SP adapter name and ID are **OpenToken SpSample** and `opentokenSpSample`, respectively. In the **SP Configuration > Server Settings > Federation Info** screen, the base URL for your PingFederate SP server is defined as `https://sso.xray.local:9071`. There are no other IdP connections besides those required to connect with Alpha and Bravo.

This example requires the following components:

- An SP adapter instance deployed, configured, and integrated with the target application.
- An IdP connection to the partner (*step 1*).
- An IdP connection to the third-party IdP that facilitates the multifactor authentication process (*step 2*).
- An authentication policy contract to carry user attributes from the partner to the target application (*step 3*).
- An SP authentication policy (*step 4* and *step 7*).
- An adapter mapping between the authentication policy contract and the applicable SP adapter instance (*step 5*).
- An SP-initiated SSO URL (*step 6*).

To fulfill the requirements:

1. In the **SP Configuration** menu, click **Create New** (under **IdP Connections**).
 - a) Follow the connection wizard to create a SAML 2.0 IdP connection to Alpha; in this example Alpha's entity ID is `sso.alpha.local`.
 - b) In the **IdP Connection > Browser SSO > SAML Profiles** screen, make sure that the **SP-Initiated SSO** check box is selected.
 - c) In the **IdP Connection > Browser SSO > User-Session Creation > Identity Mapping** screen, 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 (**OpenToken SpSample**) to be mapped into the connection in the **IdP Connection > Browser SSO > User-Session Creation > Target Session Mapping** screen. The rest of the steps remain unchanged.

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 **SP Configuration > Policies** screen. In the **Authentication Policy Contracts** screen, click **Create New Contract**.

 **Tip:** 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).

a) Follow the wizard to create an authentication policy contract.

b) In the **Contract Attributes** screen, extend the contract two attributes, subject and mail; for example:




c) Complete the rest of the connection configuration.

4. Go to the **SP Configuration > Policies** screen to define a policy to enforce third-party authentication.

a) Select **sso.alpha.local (IdP Connection)** from the **Action** list.

In general, an IdP connection has two results, **Fail** and **Success**; for example:




Each result forms its own path that requires further configuration

b) For the **sso.alpha.local (IdP connection) > Fail** path, leave its action as **--Done--**.

The policy engine stops. PingFederate rejects the SSO request.

c) For the **sso.alpha.local (IdP connection) > Success** path, select **sso.bravo.local (IdP Connection)** from its **Action** list (which also forms two paths); for example:



- d) Click **Options** (underneath **sso.bravo.local (IdP Connection)**) to relay the user ID (SAML_SUBJECT) from **sso.alpha.local** to **sso.bravo.local**.

In the **Incoming User ID** dialog, select **IdP Connection (sso.alpha.local)** from the **Source** list and **SAML_SUBJECT** from the **Attribute** list; for example:

Click **Done**.

- e) For the **sso.bravo.local (IdP connection)** > **Fail** path, leave its action as **--Done--**.

The policy engine stops. PingFederate rejects the SSO request.

- f) For the **sso.bravo.local (IdP connection)** > **Success** path, select the authentication policy contract created in [step 3](#); for example:

- g) Click **Contract Mapping** (underneath the authentication policy contract) and follow the **Manage Authentication Policies > Authentication Policy Contract Mapping** wizard to configure the fulfillment of the authentication policy contract.

Optional: In the **Attribute Sources & User Lookup** screen, click **Add Attribute Source** to configure data store queries.

In the **Contract Fulfillment** screen, select **IdP Connection (sso.alpha.local)** from the **Source** list and the appropriate attributes in the **Value** list to fulfill the authentication policy contract attributes; for example:

Contract Fulfillment	Source	Value	Actions
mail	IdP Connection (sso.alpha.local)	mailEmail	None available
idpuid	IdP Connection (sso.alpha.local)	SAML_SUBJECT	None available



Tip: Essentially, you are mapping attributes to an authentication policy contract from the authentication policy.

Optional: In the **Issuance Criteria** screen, configure conditions to be validated before issuing an authentication policy contract. (For more information, see [Define issuance criteria for authentication policy contract mapping](#) on page 222.)

Click **Done** and then **Save**.


5. Go to **SP Configuration > Adapter Mappings** screen. In the **Authentication Policy Adapter Mappings** screen, map the authentication policy contract to the SP adapter instance that you have deployed, configured, and integrated with the target application.
 - a) Select the authentication policy contract (created in [step 3](#)) from the **Source Instance** list.
 - b) Select the applicable SP adapter instance from the **Target Instance** list.
 - c) Click **Add Mapping**.
 - d) Follow the **Authentication Policy Adapter Mappings > Mapping Configuration** wizard to create the mapping.

Optional: In the **Attribute Sources & User Lookup** screen, click **Add Attribute Source** to configure data store queries.

Optional: In the **Target App Info** screen, enter application information.

In the **Adapter Contract Fulfillment** screen, select **Authentication Policy Contract** from the **Source** list and the appropriate contract attributes in the **Value** list to fulfill the SP adapter contract; for example:

SP Adapter Contract	Source	Value	Action
spAdapterDefault	Authentication Policy Contract	spEmail	None available
spAdapterDefault	Authentication Policy Contract	subject	None available

 **Tip:** Basically, you are mapping attribute values from the authentication policy contract to the target application through the applicable SP adapter instance.

Optional: In the **Default Target URL** screen, specify a default target URL for this mapping configuration.

Optional: In the **Issuance Criteria** screen, configure conditions to be validated before issuing an SP adapter contract. (For more information, see [Define issuance criteria for authentication policy contract mapping](#) on page 222.)


Click **Done**.

- Configure an SP-initiated SSO URL in your target application by combining the base URL of your PingFederate SP server (`https://sso.xray.local:9071`), 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:9071/sp/startSSO.ping?
SpSessionAuthnAdapterId=opentokenSpSample
```

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 the parameter `TargetResource` (or `TARGET`) is used and includes its own query parameters, the parameter value must be URL-encoded.

Any other parameters that contain restricted characters (many SAML URNs, for example) also must be URL-encoded.

For information about URL encoding, please refer to third party resources, such as [HTML URL-encoding Reference](#) (www.w3schools.com/tags/ref_urlencode.asp).

- When you are ready to enable SP authentication policies, go to the **SP Configuration > Policies** screen, select the **Enable SP Authentication Policies** check box, and click **Save**.


Configure SP authentication policies for users from multiple IdPs

You can configure SP authentication policies to handle different authentication requirements for multiple IdP connections.

Suppose you configured the following use cases in PingFederate® 8.0:

- Two SP adapter instances in the **SP Configuration > Adapters** screen:

Instance Name	Instance Id	Type	Parent Name	Action
Sample	sample	Opentoken Adapter 2.0.0	Sample	Details Check Usage
Sample Delta	sampleDelta	Opentoken Adapter 2.0.0	Sample	Details Check Usage

 **Tip:** In this example, because the applications use the same SP adapter type, the same configuration, and the same set of user attributes (subject and spEmail), **Sample Delta** is a child adapter of **Sample** without any overridden settings.

- Three entries in the **SP Configuration > Target URL Mapping** screen:

URL	Target Type	Target Session	Action
https://sso.xray.local:9031/sp/sample/MainPage?pp=Alpha*	SP Adapter	Sample	Get Details
https://sso.xray.local:9031/sp/sample/MainPage?pp=Charlie*	SP Adapter	Sample	Get Details
https://sso.xray.local:9031/sp/sample/MainPage?pp=Delta*	SP Adapter	Sample Delta	Get Details

- Three IdP connections to your partners as follows:

Partner (Federation ID)	Identity Mapping	Attribute Contract	Target Session Mapping SP adapter instance name (SP adapter instance ID)
Alpha (sso.alpha.local)	Account Mapping	SAML_SUBJECT email	Sample (sample)
Charlie (sso.charlie.local)	Account Mapping	SAML_SUBJECT email	Sample (sample)
Delta (sso.delta.local)	Account Mapping	SAML_SUBJECT email	Sample Delta (sampleDelta)

Note that all partners support SAML 2.0 and only the SP-initiated SSO profile.

- SSO URLs for users from Alpha, Charlie, and Delta, as follows:

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%2F%3Fapp%3DAlpha%26tab%3Dhome
Charlie	https://sso.xray.local:9031/sp/startSSO.ping? PartnerIdpId=sso.charlie.local&TargetResource=https%3A%2F%2Fsso.xray.local%3A9031%2FSpSample%2FMainPage%2F%3Fapp%3DCharlie%26tab%3Dhome
Delta	https://sso.xray.local:9031/sp/startSSO.ping? PartnerIdpId=sso.delta.local&TargetResource=https%3A%2F%2Fsso.xray.local%3A9031%2FSpSample%2FMainPage%2F%3Fapp%3DDelta%26tab%3Dhome

After upgrading to PingFederate 8.1, you have the following new requirements:

- Create new IdP connections to three new partners: Echo, Foxtrot and Golf.
- Enforce multifactor authentication for users from Alpha, Charlie, Echo, and Golf through Bravo, which requires a user ID to be passed in from the original source and returns only the user ID when the users fulfill the multifactor authentication requirement.

The new required components are:

- Two additional SP adapter instances (*step 1*):
 - **Sample Echo** to integrate with Echo's target application.
 - **Sample Golf** to integrate with Golf's target application.
- Four new IdP connections (*step 2*, *step 3*, and *step 4*):

Partner (Federation ID)	Identity Mapping	Attribute Contract	Target Session Mapping SP adapter instance name (SP adapter instance ID)
Bravo (sso.bravo.local)	No Mapping	SAML_SUBJECT (without other attributes)	N/A
Echo (sso.echo.local)	No Mapping	SAML_SUBJECT email	N/A
Foxtrot (sso.foxtrot.local)	Account Mapping	SAML_SUBJECT email	Sample (sample)
Golf (sso.golf.local)	No Mapping	SAML_SUBJECT email	N/A

Note that all partners support SAML 2.0 and only the SP-initiated SSO profile.

- Three authentication policy contracts (*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 to determine if a request is meant for Alpha or Charlie, because Alpha's and Charlie's target applications share an SP adapter instance (*step 6*).
- Three SP authentication policies to enforce the multifactor authentication requirement (*step 7*, *step 8*, and *step 12*).
- Three adapter mappings between the authentication policy contracts and the applicable SP adapter instances (*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 (*step 10*):
- SSO URLs for all partners (*step 11*).

Follow these steps to fulfill the new requirements:

1. Go to the **SP Configuration > Adapters** screen to create two new SP adapter instance: **Sample Echo** and **Sample Golf**; for example:


Instance Name	Instance ID	Type	Parent Name	Action
Sample	sample	OpenSAML Adapter 2.0.0		Delete Check Usage
Sample Echo	sampleEcho	OpenSAML Adapter 2.0.0	Sample	Delete Check Usage
Sample Golf	sampleGolf	OpenSAML Adapter 2.0.0	Sample	Delete Check Usage




Tip: In this example, because the applications use the same SP adapter type, the same configuration, and the same set of user attributes, **Sample Echo** and **Sample Golf** are child adapters of **Sample** without any overridden settings.


2. Create an IdP connection to Bravo.
 - a) In the **SP Configuration** menu, click **Create New** (under **IdP Connections**) and follow the **IdP Connection** wizard to create a SAML 2.0 IdP connection.
 - b) In the **IdP Connection > Browser SSO > SAML Profiles** screen, select the **SP-Initiated SSO** check box.


- c) In the **IdP Connection > Browser SSO > User-Session Creation > Identity Mapping** screen, select **No Mapping**.
 - d) Complete the rest of the connection configuration.
3. Create IdP connections to Echo and Golf.
 - a) In the **SP Configuration** menu, click **Create New** (under **IdP Connections**) and follow the **IdP Connection** wizard to create a SAML 2.0 IdP connection to Echo (and then to Golf).
 - b) In the **IdP Connection > Browser SSO > SAML Profiles** screen, select the **SP-Initiated SSO** check box.

 **Tip:** If the partner and your organization agree to support other profiles, select them accordingly.
 - c) In the **IdP Connection > Browser SSO > User-Session Creation > Identity Mapping** screen, 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 Echo** or **Sample Golf**) to be mapped into the connection in the **IdP Connection > Browser SSO > User-Session Creation > Target Session Mapping** screen. The rest of the steps remain unchanged.
 - d) In the **IdP Connection > Browser SSO > User-Session Creation > Attribute Contract** screen, extend the contract with the email attribute.
 - 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) In the **SP Configuration** menu, click **Create New** (under **IdP Connections**) and follow the **IdP Connection** wizard to create a SAML 2.0 IdP connection.
 - b) In the **IdP Connection > Browser SSO > SAML Profiles** screen, select the **SP-Initiated SSO** check box.

 **Tip:** If the partner and your organization agree to support other profiles, select them accordingly.
 - c) In the **IdP Connection > Browser SSO > User-Session Creation > Identity Mapping** screen, select **Account Mapping**.
 - d) In the **IdP Connection > Browser SSO > User-Session Creation > Attribute Contract** screen, extend the contract with the email attribute.
 - e) In the **IdP Connection > Browser SSO > User-Session Creation > Target Session Mapping** screen, click **Map New Adapter Instance** and follow the **Adapter Mapping & User Lookup** wizard to map the attributes from the assertion to the SP adapter instance **Sample**.
 - f) Complete the rest of the connection configuration.
 5. Create three authentication policy contracts.

 **Tip:** 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 7* in this example).
 - You map attributes from authentication policy contracts to target applications through adapter mappings (*step 9* in this example).
 - a) Go to the **SP Configuration > Policy Contracts** screen, click **Create New Contract** and follow the **Authentication Policy Contract** wizard to create an authentication contract for users from Alpha and Charlie (Echo, and then Golf).
 - b) In the **Contract Attributes** screen, extend the authentication policy contract with an attribute for user's email address; for example, apcEmail.
 - c) Complete the rest of the connection configuration.
 - d) Repeat these steps to create an authentication policy contract for users from Echo, and then for users from Golf; for example:

Contract Name	Contract ID	Action
Authenticator	Z096qPIdpId04	Direct
Echo authenticated	Y020PIdpId04	Direct
Golf authenticated	W0496666604	Direct


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 Authentication Selector to categorize requests based on their respective PartnerIdpId query parameter values at runtime. Based on the results, the policy diverts accordingly.

Create an instance of the HTTP Request Parameter Authentication Selector.


- Go to the **SP Configuration > Selectors** screen and click **Create New Instance**.
- In the **Type** screen, select **HTTP Request Parameter Authentication Selector** from the **Type** list and provide a name and ID for the selector instance.
- In the **Authentication Selector** screen, enter PartnerIdpId in the **HTTP Request Parameter Name** field.
- In the **Selector Result Values** screen, 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.


- Complete the rest of the configuration; for example:



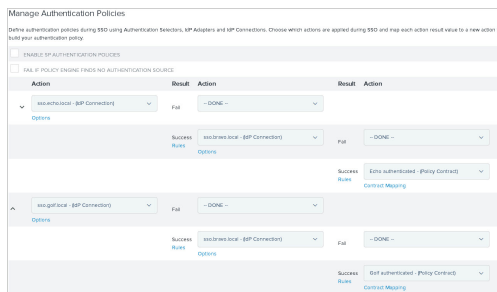
7. Go to the **SP Configuration > Policies** screen to define a new authentication policy to enforce third-party authentication for users from Echo (and then for users from Golf).

 **Tip:** If you need more information about each sub step, see [step 4](#) in [Configure an SP authentication policy for users from one IdP](#) on page 224.

- Select **sso.echo.local (IdP Connection)** from the **Action** list.
- For the **sso.echo.local (IdP connection) > Fail** path, leave its action as **--Done--**.
- For the **sso.echo.local (IdP connection) > Success** path, select **sso.bravo.local (IdP Connection)** from its **Action** list.
- Click **Options** (underneath **sso.bravo.local (IdP Connection)**) to relay the user ID (SAML_SUBJECT) from `sso.echo.local` to `sso.bravo.local`.
- For the **sso.bravo.local (IdP connection) > Fail** path, leave its action as **--Done--**.
- For the **sso.bravo.local (IdP connection) > Success** path, select the authentication policy contract **Echo authenticated** from its **Action** list, and then click **Contract Mapping** (underneath the authentication policy contract) to configure the fulfillment of the authentication policy contract.

 **Tip:** Essentially, you are mapping attributes to an authentication policy contract from the authentication policy.

- Repeat these steps to define a new authentication policy to enforce third-party authentication for users from Golf; for example:



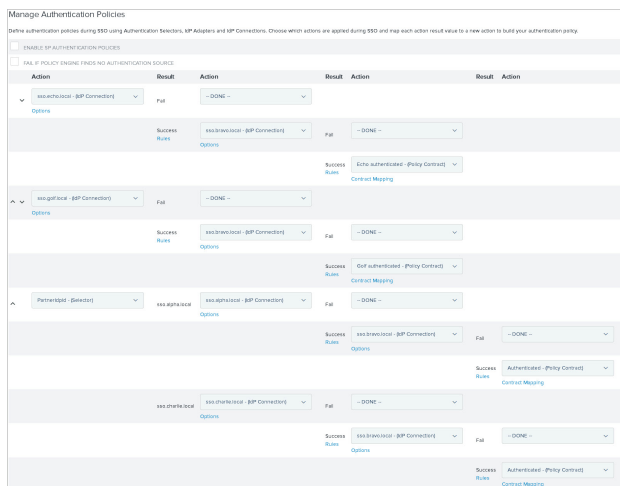
8. In the **Manage Authentication Policies** screen, define a new authentication policy to enforce third-party authentication for users from Alpha and Charlie.



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 policy paths that are meant for the respective IdPs. This example uses an instance of the HTTP Request Parameter Authentication Selector created in [step 6](#).

- a) Select the instance of the HTTP Request Parameter Authentication Selector from the **Action** list.
- b) For the path with result of `sso.alpha.local`, repeat [step 7](#) to configure the authentication requirement for the `sso.alpha.local` IdP connection.
- c) For the path with result of `sso.charlie.local`, repeat [step 7](#) to configure the authentication requirement for the `sso.charlie.local` IdP connection.

At this point, the SP authentication policies are configured as follows:




d) Click **Save**.

9. Create adapter mappings.

- a) Go to the **SP Configuration > Adapter Mappings** screen.
- b) Create three adapter mappings to map from the authentication policy contracts (created in [step 5](#)) to the applicable SP adapter instances (**Sample**, **Sample Echo**, and **Sample Golf**); for example:

Authentication Policy Adapter Mappings	
Authentication to Sample	Action
Echo authenticated to Sample Echo	Action
Golf authenticated to Sample Golf	Action

For each adapter mapping:




Tip: Basically, you are 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 the **SP Configuration > Target URL Mapping** screen and add the following mappings:

URL	Target Session
<code>https://sso.xray.local:9031/SpSample/MainPage/?app=Echo&*</code>	Sample Echo
<code>https://sso.xray.local:9031/SpSample/MainPage/?app=Foxtrot&*</code>	Sample
<code>https://sso.xray.local:9031/SpSample/MainPage/?app=Golf&*</code>	Sample Golf

At this point, the target URL mappings are configured as follows:



11. Configure the following SSO URLs:

Partner	SSO URL
Alpha	<p><code>https://sso.xray.local:9031/sp/startSSO.ping?PartnerIdpId=sso.alpha.local&TargetResource=https%3A%2F%2Fsso.xray.local%3A9031%2FSpSample%2FMainPage%2F%3Fapp%3DAlpha%26tab%3Dhome</code></p> <p>The SSO URL has not changed.</p> <p>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.</p>
Charlie	<p><code>https://sso.xray.local:9031/sp/startSSO.ping?PartnerIdpId=sso.charlie.local&TargetResource=https%3A%2F%2Fsso.xray.local%3A9031%2FSpSample%2FMainPage%2F%3Fapp%3DCharlie%26tab%3Dhome</code></p> <p>The SSO URL has not changed.</p> <p>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.</p>
Delta	<p><code>https://sso.xray.local:9031/sp/startSSO.ping?PartnerIdpId=sso.delta.local&TargetResource=https%3A%2F%2Fsso.xray.local%3A9031%2FSpSample%2FMainPage%2F%3Fapp%3DDelta%26tab%3Dhome</code></p>

Partner	SSO URL
	<p>The SSO URL has not changed.</p> <p>Optional: Based on the current configuration, you could remove the PartnerId query parameter because it is not required. You could also replace the TargetResource query parameter and its value with the SpSessionAuthnAdapterId query parameter and the applicable SP adapter instance ID (SpSessionAuthnAdapterId=sampleDelta) if you have configured a default target URL in the IdP connection to Delta (or a default SP SSO URL for all IdP connections).</p>
Echo	<pre>https://sso.xray.local:9031/sp/startSSO.ping? SpSessionAuthnAdapterId=sampleEcho&TargetResource=https%3A%2F %2Fsso.xray.local%3A9031%2FSpSample%2FMainPage%2F%3Fapp%3DEcho %26tab%3Dhome</pre> <p>This is a new SSO URL.</p> <p>Optional: Based on the current configuration, you could remove the TargetResource 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).</p>
Foxtrot	<pre>https://sso.xray.local:9031/sp/startSSO.ping? SpSessionAuthnAdapterId=sample&TargetResource=https%3A%2F %2Fsso.xray.local%3A9031%2FSpSample%2FMainPage%2F%3Fapp %3DFoxtrot%26tab%3Dhome</pre> <p>This is a new SSO URL.</p> <p>Optional: Based on the current configuration, you could remove the TargetResource query parameter and its value if you have configured a default target URL in the IdP connection to Foxtrot (or a default SP SSO URL for all IdP connections).</p>
Golf	<pre>https://sso.xray.local:9031/sp/startSSO.ping? SpSessionAuthnAdapterId=sampleGolf&TargetResource=https%3A%2F %2Fsso.xray.local%3A9031%2FSpSample%2FMainPage%2F%3Fapp%3DGolf %26tab%3Dhome</pre> <p>This is a new SSO URL.</p> <p>Optional: Based on the current configuration, you could remove the TargetResource query parameter and its value if you have configured a default target URL in the IdP connection to Golf (or a default SP SSO URL for all IdP connections).</p>



Important: When the parameter TargetResource (or TARGET) is used and includes its own query parameters, the parameter value must be URL-encoded.

Any other parameters that contain restricted characters (many SAML URNs, for example) also must be URL-encoded.

For information about URL encoding, please refer to third party resources, such as [HTML URL-encoding Reference](http://www.w3schools.com/tags/ref_urlencode.asp) (www.w3schools.com/tags/ref_urlencode.asp).

- When you are ready to enable SP authentication policies, go to the **SP Configuration > Policies** screen, select the **Enable SP Authentication Policies** check box, and click **Save**.

Configure 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 an SP and an IdP connection on the same server.

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 example, you configured the following use cases in PingFederate® 8.0.

For users from Echo (an IdP):

- An SP adapter instance:
 - Name: Sample Echo
 - ID: sampleEcho
- A WS-Federation IdP connection:
 - Partner: Echo
 - Federation ID: sso.echo.local
 - Identity mapping method: Account mapping
 - Default target URL: `https://sso.xray.local:9031/SpSample/MainPage/?app=Echo&tab=home`
 - SSO URL: `https://sso.xray.local:9031/sp/startSSO.ping?PartnerIdpId=sso.echo.local`

For internal users:

- An IdP HTML Form Adapter instance (HTML Form) validating credentials through a Password Credential Validator instance against your user directory
- An adapter-to-adapter mapping:
 - Source: HTML Form
 - Target: Sample Echo
 - Default target URL: `https://sso.xray.local:9031/SpSample/MainPage/?app=Internal&tab=home`
 - SSO URL: `https://sso.xray.local:9031/pf/adapter2adapter.ping?SpSessionAuthn=sampleEcho`

After upgrading to PingFederate 8.1, if you want to enforce multifactor authentication for users from Echo through Bravo, you can create an IdP connection to Bravo and the following authentication policy:

Action	Result	Action	Result	Action
sso.bravo.local - SP Connectors	Fail	--DONE--		
sso.bravo.local - SP Connectors	Fail	--DONE--		
Echo authenticated - Policy Contract	Success	Contract Mapping		

Because the authentication policy ends with a policy contract **Echo authenticated**, you must create an adapter mapping (from the policy contract) to **Sample Echo**, the SP adapter instance integrated with the target application. You also need to update the SSO URL for users from Echo to:

```
https://sso.xray.local:9031/sp/startSSO.ping?SpSessionAuthnAdapterId=sampleEcho
```

When you select the **Enabled SP Authentication Policies** check box to turn on the policy, the `/pf/adapter2adapter.ping` endpoint is disabled to prevent malicious Echo'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, thus circumventing 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 (*step 1*).

- An authentication policy contract, **Internal authenticated**, to carry attributes from internal users to the target applications. (*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 (*step 3*).
- An authentication policy for the internal users (*step 4*).
- An adapter mapping to map from the authentication policy contracts **Internal authenticated** to the SP adapter instance **Sample Internal** (*step 5*)
- A new SSO URL for the internal users (*step 6*).

Follow these steps to fulfill the new requirements:

1. Go to the **SP Configuration > Adapters** screen to create a new SP adapter instance; for example, **Sample Internal** (with an SP adapter ID of `sampleInternal`).



Tip: In this example, because the applications use the same SP adapter type, the same configuration, and the same set of user attributes (subject and spEmail), **Sample Internal** is a child adapter of **Sample Echo** without any overridden settings.

2. Go to the **SP Configuration > Policy Contracts** screen to create a new authentication policy contract; for example, **Internal authenticated**.



Tip: 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).

3. Go to the **SP Configuration > Selectors** screen to create an instance of the CIDR Authentication Selector with one or more network ranges that correspond to your internal users.



Note: 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. Another approach is to 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.

4. Go to the **SP Configuration > Policies** screen to configure your policies as follows.

Action	Result	Action	Result	Action
Authentication (SP Connector)	Fail	--NONE--		
Success Rule	Authentication (SP Connector)	Fail	--NONE--	
Success Rule	Internal Authenticated (Policy Contract)	Fail	--NONE--	
HTML Form (Adapter)	Fail	--NONE--		
Success Rule	PingID (Adapter)	Fail	--NONE--	
Success Rule	Internal Authenticated (Policy Contract)	Fail	--NONE--	

Figure 19: With PingID Adapter

Figure 20: With CIDR Authentication Selector

- Go to the **SP Configuration > Adapter Mappings** screen to create a mapping from the new authentication policy contract (Internal authenticated) to the new SP adapter instance (sampleInternal).

Enter `https://sso.xray.local:9031/SpSample/MainPage/?app=Internal&tab=home` in the **Default Target URL** screen.

- 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, provide PingFederate administrators the following benefits:

- The capability to build an attribute contract with attribute values from multiple authentication sources or data store queries through an authentication policy.
- The flexibility to map only the policy contract to a connection. Authentication sources in the policy leading up to the contract are not required to be mapped into the connection. 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 SP connections, 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 (see [Federation hub use cases](#) on page 86).

Manage policy contracts

You manage authentication policy contracts (formerly known as connection mapping contracts) in the **Authentication Policies > Policy Contracts** screen.

- Click **Create New Contract** to create a new authentication policy contract.
- Edit an existing authentication policy contract by clicking its name.
- Check its usage.
- Use the **Delete** or **Undelete** workflow to remove an authentication policy contract (if the authentication policy contract is not in use) or cancel the removal request.

Edit contract information

You can enter or modify the contract name in the **Contract Info** screen. When complete, click **Done**.

Define contract attributes

You manage the user attributes in the authentication policy contract in the **Contract Attributes** screen. Every authentication policy contract comes with a `subject` attribute. You extend the contract with additional attributes, as needed.

- Enter an attribute under the **Extended the Contract** column, and then click **Add**.

Attribute names are case-sensitive and must suit the needs of your partners. Repeat to add more attributes as needed.

- Modify an existing attribute or undo changes made by using the **Edit**, **Update**, or **Cancel** workflow.
- Click **Delete** to remove an existing attribute.

Review the policy contract

When you have finished configuring authentication policy contract, you can review the configuration on the **Summary** screen.

- If you need to make any changes, click the heading over the information you want to edit. When you are finished, click **Done** and then **Save**.

Adapter Mappings

This configuration allows attributes from an authentication policy contract to be mapped directly to an SP adapter instance, allowing 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 path, and to apply the authentication policy contract to the target application.

Configure authentication policy adapter mappings

1. Go to **SP Configuration > Adapter Mappings** screen.
2. Select the applicable authentication policy contract from the **Source Instance** list.
3. Select the SP adapter instance integrated with your target application from the **Target Instance** list.
4. Click **Add Mapping**.
5. Follow the **Authentication Policy Adapter Mappings > Mapping Configuration** wizard to create the mapping.

- a) Optional: In the **Attribute Sources & User Lookup** screen, click **Add Attribute Source** to configure one or more data store queries to fulfill the SP adapter contract.

Queries are executed in the order as shown on the **Attribute Sources & User Lookup** screen. Use the up and down arrows as needed to adjust the order.

If a required attribute (such as the user identifier of an adapter) cannot be fulfilled, the request fails.

For more information, see [Fulfillment by data store queries](#) on page 506.

- b) Optional: In the **Target App Info** screen, enter application information.
- c) In the **Adapter Contract Fulfillment** screen, 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 data store queries, dynamic texts, or results from OGNL expressions.
- d) Optional: In the **Default Target URL** screen, specify a default target URL for this mapping configuration.
- e) Optional: In the **Issuance Criteria** screen, configure conditions to be validated before issuing an SP adapter contract (see [Define issuance criteria for adapter mapping](#)).
- f) In the **Summary** screen, review the configuration and modify as needed. When complete, click **Done**.

6. In the **Authentication Policy Adapter Mappings** screen, click **Save**.

Define issuance criteria for adapter mapping


Use the **Issuance Criteria** screen to define criteria that PingFederate® can evaluate to determine whether to approve or reject requests for protected resource.

The configuration begins by selecting a source containing the attribute to be verified. Some sources, such as **Context** and **Mapped Attributes**, are common to almost all use cases. Other sources, such as **Assertion**, have dependency on the type of configuration, for which the administrative console automatically hides the irrelevant sources automatically. After selecting a source, choose an attribute to be verified. Depending on the selected source, attributes vary. For example, if **Account Linking** is the selected source, **Local User ID** is the available attribute. Finally, select the comparison method, specify the desired value, and enter an optional error message for scenarios where the condition fails.

When multiple criteria are configured through the user interface, all criteria must be met at runtime (evaluated as *true*) for the request to succeed; otherwise PingFederate rejects the request and returns an error message. If only a subset or a layer of the criteria is required, use one or more OGNL expressions to define them.


1. Select a source containing the attribute to be verified.

For more information about the **Source** list, refer to the following table:


Source	Description
Authentication Policy	Select to access attributes from the authentication policy contract.
Context	Select to use values 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.
JDBC, LDAP, or Custom (if configured)	Select to access attributes returned from a data source.
Mapped Attributes	Select to access the mapped attributes.

2. Select an attribute from the **Attribute Name** list.

3. Select a comparison method from the **Condition** list.

 **Note:** Use multi-value conditions when you want PingFederate to verify that an attribute contains or does not contain the specified value in its attribute list. Using a single-value condition when an attribute has multiple values causes the condition to evaluate consistently to false.

4. Enter an exact value in the text field.

 **Note:** If your use case requires more complex evaluations, such as partial matching, use one or more OGNL expressions to define them.

5. Optional: Enter a custom error message in the **Error Result** field that PingFederate returns when the criterion fails.

If you leave this field blank, PingFederate returns `ACCESS_DENIED` at runtime when the criterion fails.


6. Click **Add**.


Use the **Edit**, **Update**, **Cancel**, **Delete**, and **Undelete** links to manage the criterion.

7. Optional: Repeat these steps to specify additional criteria.

Define issuance criteria using OGNL expressions

1. Click **Show Advanced Criteria**.

 **Note:** Expressions must be enabled for the **Show Advanced Criteria** button to appear.

 **Tip:** If you have not already done so, you may enable OGNL expression by editing the `org.sourceid.common.ExpressionManager.xml` file in the `<pf_install>/pingfederate/server/default/data/config-store` directory. 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 on the **Server Configuration > Cluster Management** screen, and restart all nodes.

2. Specify a criterion.

- a) Enter an OGNL expression into the text field.
- b) Enter a custom error message in the **Error Result** field.



Note: If the OGNL expression resolves to a string value (rather than *true* or *false*), the returned value overrides the custom error message.

- c) Click **Add**.
- d) Click the **Test** link to input values and test the resulting output for the expression.

Use the **Edit**, **Update**, **Cancel**, **Delete**, and **Undelete** links to manage the OGNL expression.

3. Optional: Repeat these steps to specify another OGNL expression.

If multiple OGNL expressions are specified, they all must pass in order for the request to succeed. Alternatively, you can combine multiple criteria in one OGNL expression.

Sessions

Authentication sessions control when a user is redirected back to the authentication system on subsequent Browser SSO requests. The global policy and timeout settings apply to all authentication sources, as described in the following table:

Authentication source	Use cases
IdP adapter instance	<ul style="list-style-type: none"> • IdP-initiated requests received at the <code>/idp/startSSO.ping</code> application endpoint • Authentication requests (AuthnRequest) received at the <code>/idp/SSO.saml2</code> protocol endpoint • OAuth authorization requests using IdP adapter instances for authentication (through the IdP Adapter Mapping configuration) • Adapter-to-adapter mappings (through the Adapter-to-Adapter Mappings configuration)
IdP connection	<ul style="list-style-type: none"> • SP-initiated requests received at the <code>/sp/startSSO.ping</code> application endpoint • OAuth authorization requests using IdP connections for authentication (through the OAuth Attribute Mapping configuration)

Administrators may enable authentication sessions for all authentication sources, with or without specifying overrides for selected authentication sources.

Alternatively, administrators may enable authentication sessions for selected authentication sources (and optionally with its own set of overrides).

The override options include:

- Disable or enable authentication session; PingFederate® always redirects the users back to the authentication system.
- Override the idle timeout, the maximum timeout, or both, in minutes.

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 period, the session is extended by the idle timeout value (another 60 minutes by default). A 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. Note that the authentication system may or may not challenge the user to complete an authentication process based on its own session management policy or processing logic.

- 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.

Administrators may also leverage the SLO application endpoints (`/idp/startSLO.ping` and `/sp/startSLO.ping`) to terminate adapter sessions. For example, when the **Track Adapter Sessions for Logout** check box is selected, if a user triggers a web application to send a logout request to one of the SLO application endpoints, the user is logged out locally (including adapter sessions established through **Adapter-to-Adapter Mappings** configuration). If the applicable connections support SLO (for instance an SLO-enabled SAML 2.0 connection or a WS-Federation connection), logout requests are sent to the partners based on the protocol specifications.

1. Go to the **Authentication Sessions** screen from the **IdP Configuration** or **SP Configuration** menu.
2. Optional: Select the **Track Adapter Sessions for Logout** check box to enable SLO for all adapter instances on a per-user basis.
3. Configure the global policy and settings.
 - a) Select the **Enable Sessions for All Authentication Sources** check box to enable authentication sessions for all authentication sources.

Clear this check box if you prefer to enable authentication sessions for only a few authentication sources.
 - b) Optional: Modify the default idle timeout value (in minutes).

The default value is 60 minutes.
 - c) Optional: Modify the default maximum timeout value.

The default value is 480 minutes (8 hours).
4. Optional: Select an authentication source from the list to manage its session settings and overrides.
 - If the **Enable Sessions for All Authentication Sources** is selected, clear the check box under **Enable Sessions** if you want to disable authentication sessions for the selected authentication source.
 - If the **Enable Sessions for All Authentication Sources** is not selected, select the check box under **Enable Sessions** if you want to enable authentication sessions for the selected authentication source.
 - To override its timeout values (from the default values), select the check box under **Override Timeouts** and then enter values into the fields under **Idle Timeout** and **Max Timeout** (in minutes).
 - To enforce authentication requirement based on the authentication context for the selected authentication source, select the **Authentication Context Sensitive** check box.

Click **Add** to complete the entry.

Repeat this step to configure overrides for another authentication source.


Use the **Edit**, **Update**, and **Cancel** workflow to make or undo a change to an existing value. Use the **Delete** and **Undelete** workflow to remove an existing value or cancel the removal request.

5. Click **Save** to complete the configuration.

OAuth configuration


To use the OAuth Authorization Server (AS), start by enabling that role for PingFederate under **Server Settings** on the Roles and Protocols screen. Then configure the OAuth AS using options available under Server Configuration on the Main Menu, as described in this section.

For more information about the OAuth AS, see [About OAuth](#) on page 61.

 **Tip:** Service providers may also add OAuth capabilities to the Browser SSO configuration for IdP connection partners, see [Configure OAuth attribute mapping](#).

Configure OAuth use cases

Use the **OAuth Settings** menu to configure various settings to support the grant types that your applications require.



1. Configure the authorization server (AS) settings in the **OAuth Settings > Authorization Server Settings** screen.
2. Define the default scope, add optional scopes, create scope groups from any optional scopes as needed in the **OAuth Settings > Scope Management** screen.
3. Create one or more access token management instances in the **OAuth Settings > Access Token Management** screen.

This is where you define the access-token's attribute contract for any given access token management instance.

4. Configure one or more mappings between the authentication sources and the persistent grant contract.

Authorization Code or Implicit

- Map attributes from an IdP adapter instance to the persistent grants **OAuth Settings > IdP Adapter Mapping** wizard.
- Map attributes from an IdP connection to the persistent grants using the **IdP Connection > Browser SSO > OAuth Attribute Mapping** screen.
- Create an authentication policy contract (APC) using the **Policy Contracts** wizard, 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 **OAuth Settings > Authentication Policy Contract Mapping** wizard.



Tip: If you are using a combination of authentication policies, APCs, and APC mappings, you can skip the **IdP Adapter Mapping** and **OAuth Attribute Mapping** configurations.

Resource Owner Password Credentials

- Map attributes from a password credential validator instance to the persistent grants **OAuth Settings > Resource Owner Credential Mapping** wizard.

Note that this is the first stage of the two-stage access token mapping process through the persistent grants.

5. Configure one or more mappings between the persistent grant contract (or the authentication sources directly) and the attribute contract of your access token management instances using the **OAuth Settings > Access Token Mapping** wizard.

Note that this is the second stage of the two-stage access token mapping process through the persistent grants.


6. For the SAML 2.0 or JWT profile for OAuth 2.0 authorization grants flow, configure a mapping in the **IdP Connection > OAuth Assertion Grant Attribute Mapping** screen.

Note that this use case exchanges a SAML assertion or a JWT for an OAuth access token.

7. Define one or more OpenID Connect policies in the **OAuth Settings > OpenID Connect Policy Management** wizard if you support OpenID Connect use cases and the **OpenID Connect** role (under OAuth 2.0) has been enabled in the **Server Configuration > Server Settings > Roles and Protocols** screen.
8. Create one or more OAuth clients in the **OAuth Settings > Client Management** screen.
9. Optional: Configure client session management settings.

Enable the OAuth AS role

To configure OAuth use cases, you must enable the Authorization Server (AS) role in PingFederate®.

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



1. Go to the **Server Configuration > Server Settings > Roles and Protocols** screen.
2. Select the **Enable OAuth 2.0 Authorization Server (AS) role** check box.
3. Optional: Select the **OpenID Connect** check box if you intend to support OpenID Connect use cases as well.




Configure authorization server settings

The **Authorization Server Settings** screen provides overall controls over the usage and behavior of the PingFederate® OAuth AS, including the policies and settings for various grant types, refresh-tokens, persistent grants, and ID tokens.

1. Go to the **OAuth Settings > Authorization Server Settings** screen.
2. Configure the PingFederate OAuth AS settings to suit your use cases.

Refer to the following table for detailed information about each field.

Field	Description
Authorization Code Timeout (Seconds)	The amount of time (in seconds) that an authorization code is considered valid.
Authorization Code Entropy (Bytes)	The length (in bytes) of the authorization code returned to the OAuth client.
Persistent Grant Lifetime (Blank for Indefinite)	(Integer) Indicates the number of days, hours, or minutes the OAuth AS stores persistent grants. Leave blank to maintain grants indefinitely.  Note: You can also set persistent grant expiration for specific OAuth clients, overriding this global configuration.
Refresh Token Length (Characters)	The number of characters the OAuth AS uses to generate the refresh token returned to the client.
Roll Refresh Token Values (Default Policy)	When selected, the OAuth AS generates a new refresh token value when a new access token is obtained.  Note: New refresh token values are not issued during the interval defined by the Minimum Interval to Roll Refresh Tokens field. Not selecting the check box means the refresh token value is used until it becomes invalid (either by manually revoking or by some other security setting that renders it invalid).
Minimum Interval to Roll Refresh Tokens (Hours)	The minimum number of hours that must pass before a new refresh token value can be issued. Provides a way to allow for rolling a refresh token value without having to send a new value on every request.
Reuse Existing Persistent Access Grants for Grant Types	If a client makes multiple requests for the same user and same (or lesser) scope, select the grant types you want the OAuth AS to reuse rather than creating a new grant for each request. Reusing an existing persistent grant imposes a limit of one grant per client. For example, in the case of refresh tokens, only one (the most recently issued) is valid, and the previously issued refresh token is invalidated. If multiple instances of the same client are used regularly by the same user, the associated check box should be cleared. When the Implicit grant type is selected, consent from the user is requested only for the first OAuth resource request associated with the grant. When the Authorization

Field	Description
	Code grant type is selected, the same is true <i>if</i> the Bypass Authorization for Previously Approved Persistent Grants check box is also selected.
Bypass Authorization for Previously Approved Persistent Grants	<p>When selected, consent from the user is requested only once. The user is not asked for authorization on subsequent requests until the access grant is revoked. This function applies only when using the Authorization Code grant type <i>and</i> when the Reuse Existing Persistent Access Grants for Grant Types check box is selected.</p> <p> Tip: On a per-client basis, user consent may be bypassed completely, overriding this setting.</p>
Allow Unidentified Clients to Request Resource Owner Password Credentials Grants	When selected, allows Resource Owners to obtain access tokens without defining a client.
Allow Unidentified Clients to Request Extension Grants	<p>When selected, allows user-initiated or client-initiated events (for example, a mobile application or scheduled task) to obtain access tokens without the client presenting a client_id or client_secret for the extension grant types, namely:</p> <ul style="list-style-type: none"> Validation grant type <code>urn:pingidentity.com:oauth2:grant_type:validate_bearer</code> SAML 2.0 Bearer Assertion grant type <code>urn:ietf:params:oauth:grant-type:saml2-bearer</code>
Persistent Grant Extended Attributes	Extend persistent grants to include additional attributes from your authentication systems.
Password Credential Validator	Required for HTTP Basic authentication if the OAuth REST Web Service is used for managing client applications, the OAuth Access Grant Management Service is used for managing access grants, or both. These service cannot be used if a validator is not provided.
The following settings require the selection of the OpenID Connect role in the Server Configuration > Server Settings > Roles & Protocol screen.	
Track User Sessions for Logout	When selected, PingFederate links the IdP adapter sessions that are used by OAuth clients and SLO-enabled (SAML 2.0 and WS-Federation) connections for end users. When an end user logs out of one session, PingFederate sends (via the browser) logout requests to close other associated sessions.
Revoke User Session on Logout	<p>When selected, PingFederate revokes the corresponding sessions on logout.</p> <p> Note: PingFederate always adds the session to the revocation list, even if a logout error occurs.</p>
Session Revocation Lifetime (Minutes)	<p>The amount of time (in minutes) until the revoked sessions are removed from the revocation list for optimal performance.</p> <p> Important: The value should match or exceed the maximum session lifetime of the relying parties.</p>


Define scopes

Use the **Scope Management** screen to define the default scope, add optional scopes, create scope groups from any optional scopes as needed.

1. Go to the **OAuth Settings > Scope Management** screen.
2. Configure scopes required to suite your use cases.

Note that the **Default Scope Description**, **Scope Description**, and **Scope Group Description** fields control the text appears when the user is prompted for authorization. If you want to localize the display text, you can enter a unique alias here, then use the same alias in language resource files. (Localization resource files can be found in the `<pf_install>/pingfederate/server/default/conf/language-packs` directory.)

For detailed information about each field, refer to the following table.

Field	Description
Default Scope Description (Required)	Description of the permissions implied when no scope values are indicated or in addition to any values. This description displays when the user is prompted for authorization.
Scope Value	<p>A value that represents access to a resource or API on the Resource Server (RS). Applicable values require coordination with a developer or someone familiar with details of the RS OAuth implementation.</p> <p>For OpenID Connect, you can change the message presented to users for their authorization by entering <code>openid</code> under Scope Value and the message you want under Scope Description. In addition, to use the OpenID Connect defined scopes, you <i>must</i> add the following scopes, along with descriptions:</p> <ul style="list-style-type: none"> • profile • email • address • phone <p>By default, scopes are available to all OAuth clients unless access is restricted on a per-client basis.</p> <p> Caution: Updating or removing a scope value in production will cause runtime errors if a client request specifies the scope using the previously defined value. In addition, errors will result if a requesting OAuth client is restricted to a scope whose value is no longer listed on this screen, unless the affected client configuration is also updated.</p>
Scope Description	A description of the scope value. This description appears when the user is prompted for authorization.
Scope Group Value	A value that represents a group of scopes (a <i>super scope</i>) that you can reference in applicable OAuth 2.0 protocol interactions. The scopes within the scope group can be downgraded (upon refresh) into application-specific scopes with fewer permissions.
Scope Group Description	A description of the scope group value. This description appears when the user is prompted for authorization.

OAuth access token management

PingFederate® supports multiple access token management (ATM) instances. This capability allows you (the administrator) 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.

When defining an ATM instance, you can customize various settings, including the 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. Similarly, you can add a resource server (RS) client to the ACL of multiple ATMs instances, so that only such RS 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 ATM ID (`access_token_manager_id`) or a resource URI (`aud`) in their requests to the PingFederate OAuth AS at the authorization endpoint (`/as/authorization.oauth2`), the token endpoint (`/as/token.oauth2`), and the introspection endpoint (`/as/introspect.oauth2`). For RS clients, you may configure on a per-client basis whether an RS client must specify the desired ATM instance in its token validation requests at runtime (see [Configure an OAuth client](#) on page 269).

At runtime, the PingFederate OAuth AS uses the following rules to determine which ATM instances it should use:

1. Limit 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** screen. For OAuth Assertion Grant requests, it is the set of instances for which a mapping is defined in the IdP connection. Furthermore, the ACL (if configured) 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 RS clients, you may 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 RS clients do not provide either the `access_token_manager_id` or `aud` parameter in their requests and the RS 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.

Manage access token management instances

Use the **OAuth Settings > Access Token Management** workflow to specify how the PingFederate® OAuth AS manages access tokens.

- 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**.

Define an access token management instance

Define your access token management instance on the **Type** screen.

1. Enter an instance name and an instance ID.
2. Choose the plug-in type of the access token management instance from the list.

Type varies depending on the plug-ins deployed on your server. For information about adding a customized plug-in, please contact a support representative via the [Ping Identity Support Center](#).

3. 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 screens.

Configure an access token management instance

This configuration varies depending on the Type of the access token manager. The following sections provide information for token plug-ins bundled with PingFederate:

- [Configure reference-token management](#) on page 247
- [Configure JSON-token management](#) on page 248

Configure reference-token management

Access tokens that use the reference-token data model provide a reference to some set of attributes. The resource server (RS) must de-reference the access tokens for the corresponding identity and security information at the OAuth authorization server (AS) that issued them.

Modify the default values as needed.

Refer to the following table for detailed information about each field.


Field	Description
Token Length (Required)	<p>The number of characters the AS uses to define the token reference. Increasing the length will enhance token security, if desired.</p> <p>The default and maximum values are 28 and 256, respectively.</p>
Token Lifetime (Required)	<p>The amount of time (in minutes) that an access token is considered valid.</p> <p>The default value is 120 (minutes).</p>
Lifetime Extension Policy	<p>Indicates whether the OAuth AS should reset the lifetime each time an access token is validated, subject to the values defined in the Maximum Token Lifetime and Lifetime Extension Threshold Percentage fields.</p> <p>The options are: no extension policy (the default), reset token lifetimes only for transient tokens (not backed by a persistent policy), or reset lifetimes for all tokens.</p>
Maximum Token Lifetime	<p>Defines an absolute maximum token lifetime (in minutes) for use with the Lifetime Extension Policy setting. An access token's lifetime can be extended but not beyond this setting. Any value, if specified, must be greater than or equal to the value specified in the Token Lifetime field.</p> <p>This optional field has no default value.</p>
Lifetime Extension Threshold Percentage (Required)	<p>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.</p> <p>This setting limits RPC overhead by suspending the calls until the set threshold is crossed. For example, if the token lifetime is 60 minutes and the threshold is 30%, the lifetime will not be extended until the remaining time is less than 18 minutes. This option could potentially reduce RPC traffic between nodes by orders of magnitude while still supporting a lifetime extension policy.</p> <p>The default value is 30 (percent).</p>
Advanced Fields	
Mode for Synchronous RPC	<p>Synchronous RPC calls occur when a node receives a verification request for a token it does not recognize and for token issuance.</p> <p>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.</p> <p>When All Nodes is selected, it waits for all recipients to respond.</p> <p>Majority of Nodes is the default selection.</p>
RPC Timeout (Required)	<p>The timeout value (in milliseconds) between cluster nodes during synchronous communication. The recommended value ranges from 100 milliseconds to 1000 (1 second).</p> <p>The default value is 500 (milliseconds).</p>

Configure JSON-token management


JSON Web Token (JWT) bearer access tokens are secure and self-contained tokens. This capability allows the target resource server (RS) to validate the access tokens locally or to send the access tokens to the token issuer (PingFederate®) for validation. The configuration provides for token security using either symmetric keys or asymmetric signing-certificate keys. Multiple entries are allowed for either signing mechanism to facilitate rollover of keys when they expire.

1. Add one or more symmetric keys, signing certificates, or both.

Click **Add a new row . . .**, enter information, and then click **Update** under **Action**.

 **Important:** The **Key ID** field values must be unique across all JSON-token management instances, including child instances.

If you have not yet created or imported your certificate into PingFederate, click **Manage Signing Certificates** and use the **Certificate Management** workflow to complete the task.

 **Note:** To use an RSA-based algorithm for 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.

2. Change or select entries for required fields and make other changes as needed.

Refer to the following table for information about each field. For detailed information about the algorithms, please refer to the *JSON Web Algorithms (JWA)* specification (tools.ietf.org/html/rfc7518).

Field	Description
Token Lifetime	The amount of time, in minutes, that an access token is considered valid.
(Required)	The default value is 120 (minutes).


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. Required if an asymmetric algorithm is selected in the JWE Algorithm list.
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; for example: <ul style="list-style-type: none"> • Symmetric algorithms (Direct Encryption with symmetric key, AES ...Key Wrap, and AES-GCM ... key encryption) • Asymmetric algorithms (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.

Field	Description
Active Symmetric Encryption Key ID	The ID of the key to use when using a symmetric encryption algorithm. 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 JWK format or a certificate. Applicable only if an asymmetric algorithm is selected from the JWE Algorithm list. Note that you can only specify an asymmetric encryption key here or the partner's JWKS endpoint in the Asymmetric Encryption JWKS URL field.
Asymmetric Encryption JWKS URL	The HTTPS URL of a JSON Web Key Set (JWKS) endpoint that provides a list of one or more public keys for encryption. Applicable only if an asymmetric algorithm is selected from the JWE Algorithm list. Note that 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	
Include Key ID Header Parameter	When selected (the default), the key ID is used in the kid header parameter for the token.
Include X.509 Thumbprint Header Parameter	When selected, the X.509 certificate thumbprint is used in the x5t header parameter for the token. This check box is not selected by default.
Default JWKS URL Cache Duration	When an asymmetric encryption JWKS URL is specified, in the event that the remote server does not contain any cache directives in its response, the OAuth AS only caches the content for 720 minutes (12 hours). Note that when this threshold is reached or if the cache directives indicate that the content has expired at runtime, the OAuth AS contacts the remote server to refresh the list of encryption keys from the partner.
Include JWE Key ID header parameter	When selected (the default), 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.
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, which 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. (Omitted if left blank.)
Scope Claim Name	The name of a JWT claim used to represent the scope of the grant. (Omitted if left blank.)
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.
Issuer Claim Value	The value of the Issuer claim (iss) in the JWT. (Omitted if left blank.)

Field	Description
	<p>Additionally, you may extend the contract of the access tokens with an attribute named <code>iss</code> on the Access Token Attribute Contract screen. When mapping attribute values from authentication sources to the access tokens issued by this ATM instance, the value you specify on the Contract Fulfillment screen overrides the value here.</p>
Audience Claim Value	<p>The value of the Audience claim (<code>aud</code>) in the JWT. (Omitted if left blank.)</p> <p>When no value is specified, PingFederate does not validate the Audience claim (<code>aud</code>) value if any is included in the access token.</p> <p>Additionally, you may extend the contract of the access tokens with an attribute named <code>aud</code> on the Access Token Attribute Contract screen. When mapping attribute values from authentication sources to the access tokens issued by this ATM instance, the value you specify on the Contract Fulfillment screen overrides the value here.</p>
JWT ID Claim Length	<p>Indicates the number of characters of the JWT ID (<code>jti</code>) claim in the JWT.</p> <p>The default value is 0, which means no claim is included.</p>
Access Grant GUID Claim Name	<p>The name of the JWT claim used to carry the persistent access grant GUID. (Omitted if left blank.)</p> <p>If the claim is present during validation, PingFederate checks the grant database to ensure the grant is still valid.</p> <p> Note: This use case requires that the RS must send the JWT bearer access tokens to PingFederate for validation.</p>
JWKS Endpoint Path	<p>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. The path, if specified, must begin with a forward slash; for example, <code>/oauth/jwks</code>.</p> <p>The resulting URL is <code>https://<pf_host>:<port>/ext/<JWKS Endpoint Path></code>. Furthermore, the path, if specified, must be unique across all plug-in instances, including any child instances.</p>
JWKS Endpoint Cache Duration	<p>Informs the clients the amount of time that they could cache the content from the JWKS endpoint path. Applicable only if the JWKS Endpoint Path field is configured.</p> <p>The default is 720 minutes (12 hours).</p>
Publish Key ID X.509 URL	<p>Indicates whether certificates will be made accessible by the key ID at <code>https://<pf_host>:<port>/ext/oauth/x509/kid?v=<id></code></p> <p>This check box is not selected by default.</p>
Publish Thumbprint X.509 URL	<p>Indicates whether certificates will be made accessible by thumbprint at <code>https://<pf_host>:<port>/ext/oauth/x509/x5t?v=<base64url encoded SHA-1 thumbprint></code></p> <p>This check box is not selected by default.</p>

Manage session validation settings

When an OAuth client presents an access token for validation, PingFederate® (the OAuth authorization server) checks the expiration and the other aspects of the access token. If all is well, the access token is considered valid; otherwise, the validation process returns an `invalid_grant` error to the client.

If PingFederate is configured to manage authentication sessions, you (the administrator) can optionally configure the access token validation process to evaluate the authentication sessions of the users before returning the validation results to the clients. Depending on the feature (or features) selected on the **Session Validation** screen, the validation process may return to the client an `invalid_grant` error if the associated authentication session has timed out (or expired), is not found, or has been revoked. Moreover, you may also configure PingFederate to extend the authentication sessions upon successful validations.

When any of the session validation features is enabled, the associated session identifier (pi.sri) becomes available through the access tokens. For internally managed reference tokens, the access token validation process returns the associated session identifier in the response if the access token is valid. For JWT-based access tokens, the session identifier is part of the access token. With the session identifier, an OAuth client may contact the Session Revocation API endpoint (`/pf-ws/rest/sessionMgmt/revokedSris`) to query the revocation status of an authentication session or to revoke an authentication session. For more information, see [Back-Channel Session Revocation](#) on page 273.

In essence, the session validation features enable you to conjoin the validity of access tokens and the authentication sessions of the users. Because each feature can be independently enabled or disabled per access token management (ATM) instance, you can fully customize unique API and web SSO experiences for your OAuth clients and thus the users.

It is worth noting that, once any of the session validation features is enabled for an ATM instance, clients will not be able to obtain an access token through that ATM instance by presenting a refresh token. Any attempt will result in an `unsupported_grant_type` error. The reason being is that such action, if allowed, defeats the purpose of tying the validity of access tokens and the authentication sessions of the users in the first place. Hence, it follows that the **Session Validation** features are most suitable for clients using the **Implicit** grant type because they do not use refresh tokens. That said, clients using the **Authorization Code** grant type can still take advantages of session validation as needed. They just will not be able to use refresh tokens to obtain access tokens through the ATM instances that have the session validation features enabled.

1. Go to the **OAuth Settings > Access Token Management** screen.
2. Select the applicable ATM instance or click **Create New Instance** to create a new ATM instance.

If you are creating a new ATM instance, provide the required information on the **Type** and **Instance Configuration** screens.

3. On the **Session Validation** screen, select the check box for each relevant feature.



Important: 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.

If authentication sessions are not enabled, you may still continue selecting features on this screen; however, access token validation may fail until authentication sessions are enabled.

For information about authentication sessions, including configuration steps and settings, see [Sessions](#) on page 240.

If this is a child instance, select the **Override Session Validation Settings** check box and make the adjustments as needed.

Refer to the following table for information about each feature. (Each feature is independent of each other.)

Feature	Description
Check for valid authentication session	<p>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 <code>invalid_grant</code> error.</p> <p>An authentication session is invalid when one of the following conditions applies:</p> <ul style="list-style-type: none"> • The authentication session has timed out based on the Idle Timeout field value on the Sessions screen.

Feature	Description
	<ul style="list-style-type: none"> The authentication session has expired based on the Max Timeout field value on the Sessions screen. The authentication session is not found; for example, the user has logged out. <p>You may also optimize the access token lifetime.</p> <ul style="list-style-type: none"> If this ATM instance issues internally managed reference tokens, match the value of the Token Lifetime (on the previous screen, Instance Configuration) to that of the Idle Timeout. Similarly, if you choose to specify a Maximum Token Lifetime value on the Instance Configuration screen, 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	<p>When selected, PingFederate verifies whether the session identifier (pi.sri) has been added to the revocation list. If the session has been revoked, PingFederate returns an <code>invalid_grant</code> error.</p> <p>An authentication session can be revoked via the front-channel or the back-channel. For more information, see Client session management on page 273.</p>
Update authentication session activity	<p>When selected, if the access token is valid, PingFederate also extends the lifetime of the authentication session by the Idle Timeout field value on the Sessions screen.</p>

Define the access token attribute contract

On the **Access Token Attribute Contract** screen, define the attribute contract for the access tokens issued by this access token management (ATM) instance. You must enter at least one attribute.

For JWT bearer access tokens, you may extend the attribute contract with the following attributes:

Attribute	Description
iss	<p>Adds the Issuer claim (iss) to the access token.</p> <p>When mapping attribute values from authentication sources to the access tokens issued by this ATM instance, the value you specify on the Contract Fulfillment screen overrides the Issuer Claim Value field value (if any) defined on the Instance Configuration screen.</p>
aud	<p>Adds the Audience claim (aud) to the access token.</p> <p>When mapping attribute values from authentication sources to the access tokens issued by this ATM instance, the value you specify on the Contract Fulfillment screen overrides the Audience Claim Value field value (if any) defined on the Instance Configuration screen.</p>
exp	<p>Extends the value of the Expire claim (exp), as defined by the Token Lifetime setting on the Instance Configuration screen, by the specified value (in seconds).</p>
The Client ID Claim Name field value, the Scope Claim Name field value, or the Access Grant GUID Claim Name field value (if any) defined on the Instance Configuration screen of this ATM instance.	<p>When mapping attribute values from authentication sources to the access tokens issued by this ATM instance, the values you specify on the Contract Fulfillment screen override the value of the client ID, the scope, or the persistent access grant GUID.</p>


- To add a new entry, enter the desired value and click **Add**.
- To modify an existing entry, use the **Edit**, **Update**, and **Cancel** workflow.

- To remove an existing entry, click **Delete**.

Manage resource URIs

An OAuth client can optionally include the requested resource in a query parameter (aud) when sending its request to the authorization endpoint on the PingFederate® OAuth AS.

Specify a list of resource URIs that PingFederate OAuth AS can use to select this access token management instance when the aud query parameter is provided.

 **Important:** The resource URIs must correspond to the resource expected by the resource server (RS).

- To add a new entry, enter the desired value and click **Add**.
- To modify an existing entry, use the **Edit**, **Update**, and **Cancel** workflow.
- To remove an existing entry, use the **Delete** and **Undelete** workflow.

Define access control

On the **Access Control** screen, you may restrict which OAuth clients are allowed to use this access token management instance.

1. Select the **Restrict Allowed Clients** check box.
2. Select a client from the **Allowed Clients** list and click **Add**.

Repeat this step to select additional clients, as needed.

To removal 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.

Review the access token management configuration


When you have finished the configuration, you can review it on the **Summary** screen.

If you need to make any changes, click the heading over the information you want to edit. When you are finished, click **Save**.


Manage IdP adapter mappings for OAuth

The **OAuth Settings > IdP Adapter Mapping** workflow allows you to map attributes based on an IdP adapter configuration into the USER_KEY and extended attributes for a persistent grant, and the USER_NAME (presented to the user for authorization permission). You may supplement values returned from the adapter with attribute look-ups from your user data source.

You may use this mapping for the OAuth Authorization Code and Implicit grant types.

 **Tip:** If you are using a combination of authentication policies, APCs, and APC mappings, you can skip the **IdP Adapter Mapping** and **OAuth Attribute Mapping** configurations.

- To create a mapping, select the source of the attributes from the list and click **Add Mapping**.
- To modify an existing mapping, select it by its name under **Mappings**.
- To remove an existing mapping or to 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. In addition, any corresponding entries defined in the **OAuth Settings > Access Token Mapping** screen will also be removed.

Configure IdP adapter attribute sources and user lookup

Attribute sources are specific data store or directory locations containing information that may be needed to fulfill the USER_KEY and any extended attributes of the persistent grants. In addition, attribute sources can be used to fulfill the USER_NAME presented to the end user for authorization permission. This configuration is optional.

1. Click **Add Attribute Source** (or click **Next** if you are not using a data store).
2. On the **Data Store** screen, enter values into the **Attribute Source ID** and **Attribute Source Description** fields, select an data store from the list, and click **Next** to continue the configuration.

Define grant contract fulfillment for IdP adapter mapping

On the **Contract Fulfillment** screen, you map values into the USER_KEY and extended attributes for a persistent grant, and the USER_NAME for the user's display name on the authorization page.



Note: The USER_KEY values must be unique across all end users because it is the identifier to store and to retrieve persistent grants. For example, if you have two Active Directory domains, the sAMAccountName value of an end user in one domain may collide with that of another end user in the other domain. In this scenario, you can map Subject DN to the USER_KEY.

Map each attribute from one of these Sources:

- Adapter

When you make this selection, 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: The HTTP Request 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. (If the **Expression** selection is not available, you may enable it by editing the `org.sourceid.common.ExpressionManager.xml` file in the `<pf_install>/pingfederate/server/default/data/config-store` directory.)

- LDAP/JDBC/Custom (when a data store is used)

Values are returned from your data store (if used). When you make this selection, the Value list is populated by the LDAP, JDBC or Custom attributes identified for this data store.

- Expression (when enabled)

This option provides more complex mapping capabilities—for example, transforming incoming values into different formats. All of the variables available for text entries are also available for expressions.

- Text

The value is what you enter. This can be text only, or you can mix text with references to the attributes returned from the adapter instance, using the syntax:

```
#{attribute}
```

You can also enter values from your data store, when applicable, using this syntax:

```
#{ds.attribute}
```

where `attribute` is any of the data store attributes you have selected.

1. Choose a source and then choose (or enter) a value for each attribute in the contract.
2. Click **Next**.

Specify issuance criteria for OAuth IdP adapter mapping


Use the **Issuance Criteria** screen to define criteria PingFederate® can evaluate to determine whether to issue an access token for a user. This token authorization can be used to restrict who can access an OAuth-protected resource.

To configure issuance criteria:

1. Select a source containing attributes for token authorization.

Associated attributes appear in the **Attribute Name** list:


- Adapter – Select to access attributes from the IdP Adapter.
- Context – Select to use values 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.

- LDAP/JDBC/Custom (if a data store is configured) – Select to access attributes returned from a data source.
- Mapped Attributes – Select to access the output of the mapping into the persistent grant—for example, the USER_KEY or an extended attribute.


2. Select an attribute name.

3. Select the condition you want to apply from the list.

 **Note:** Use one of the *multi-value* conditions when you want PingFederate to verify that an attribute contains or does not contain the specified value in its attribute list. Using a single-value condition when an attribute has multiple values causes the condition to evaluate consistently to false.

4. Enter the exact value in the **Value** field.

Wildcards (for partial matching) are not allowed.

 **Tip:** You can use OGNL expressions to perform more complex evaluations, including partial matching.


5. Optional: Use the **Error Result** field to customize an error code or error message for use if the token authorization fails.

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.


If you leave this field blank, the user sees a default ACCESS_DENIED error result at runtime if the authorization fails.

6. Click **Add**.


7. Repeat the steps above to add additional attributes, as needed, for each authorization criteria.

 **Note:** All criteria must pass in order for a user to be authorized.

8. Optional: Click **Show Advanced Criteria** to configure more complex evaluations using OGNL expressions.

 **Note:** Expressions must be enabled for the **Show Advanced Criteria** button to appear.

- Use the in-line editor to enter the OGNL expression.
- Use the **Error Result** field to enter an error message or an error code for use if authorization fails (see step 5 for more information).

 **Note:** If the OGNL expression resolves to a string value (rather than true or false), then the returned value overrides the Error Result.

- Click **Add**. Continue to add expressions, as needed.
- Click the **Test** link to input values and test the resulting output for the expression.

9. Click **Next** (or **Done**, if you are finished with the configuration, and then **Save** on the **Manage Mappings** screen).

To edit issuance criteria:

- Click **Edit** under Actions for the criteria.

To delete issuance criteria:

- Click **Delete** under Actions for the criteria and then click **Save**.


Review the IdP adapter mapping

When you have finished the configuration, you can review it on the **Summary** screen.

If you need to make any changes, click the heading over the information you want to edit. When you are finished, click **Done** and then **Save** on the **Manage Mappings** screen.


Manage authentication policy contract mappings

Use the **OAuth Settings > Authentication Policy Contract Mapping** screen to manage the mappings between authentication policy contracts (APCs) and the persistent grants. APCs mapped on this screen can be used to authenticate resource owners for OAuth Authorization Code and Implicit grant types.

 **Tip:** If you are using a combination of authentication policies, APCs, and APC mappings, you can skip the **IdP Adapter Mapping** and **OAuth Attribute Mapping** configurations.


1. On the **Authentication Policy Contract Mappings** screen, you may map from APCs into the persistent grants, modify a mapping as your OAuth use cases evolve, or remove a mapping when it is no longer applicable.

- To create a mapping, select the source of the attributes from the list and click **Add Mapping**.
- To modify an existing mapping, select it by its name under **Mappings**.
- To remove an existing mapping or to 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. In addition, any corresponding entries defined in the **OAuth Settings > Access Token Mapping** screen will also be removed.

2. Optional: On the **Attribute Sources & User Lookup** screen, click **Add Attribute Source** to fulfill the persistent grant contract by data store queries.


3. On the **Contract Fulfillment** screen, configure fulfillment of each attribute in the persistent grant contract.

 **Note:** The `USER_KEY` values must be unique across all end users because it is the identifier to store and to retrieve persistent grants. For example, if you have two Active Directory domains, the `sAMAccountName` value of an end user in one domain may collide with that of another end user in the other domain. In this scenario, you can map Subject DN to the `USER_KEY`.

Map each attribute from one of these Sources:

- Context

Values are returned from the context of the transaction at runtime.

 **Note:** The HTTP Request 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. (If the **Expression** selection is not available, you may enable it by editing the `org.sourceid.common.ExpressionManager.xml` file in the `<pf_install>/pingfederate/server/default/data/config-store` directory.)

- Authentication Policy Contract

When you make this selection, the associated **Value** list consists of the attributes (for example, subject) associated with the APC.

- LDAP/JDBC/Custom (when a data store is used)

Values are returned from your data store (if used). When you make this selection, the Value list is populated by the LDAP, JDBC or Custom attributes identified for this data store.

- Expression (when enabled)

This option provides more complex mapping capabilities; for example, transforming incoming values into different formats. All of the variables available for text entries are also available for expressions.

- 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 syntax:

```
${attribute}
```

You can also enter values from your data store, when applicable, using this syntax:

```
${ds.attribute}
```

where `attribute` is any of the data store attributes you have selected.

4. On the **Issuance Criteria** screen, define criteria PingFederate® can evaluate to determine whether to issue an access token for a user. This token authorization can be used to restrict who can access an OAuth-protected resource.

- a) Select a source containing attributes for token authorization.

The available choices are similar to those found in the **Contract Fulfillment** screen. Once selected, the associated attributes appear in the **Attribute Name** list.

- b) Select an attribute name.
- c) Select the condition you want to apply from the list.



Note: Use one of the *multi-value* conditions when you want PingFederate to verify that an attribute contains or does not contain the specified value in its attribute list. Using a single-value condition when an attribute has multiple values causes the condition to evaluate consistently to false.

- d) Enter an exact value for the attribute.



Tip: You can use OGNL expressions to perform more complex evaluations, including partial matching.

- e) Use the **Error Result** field to customize an error code or error message for use if the token authorization fails.

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.

If you leave this field blank, the user sees a default ACCESS_DENIED error result at runtime if the authorization fails.

- f) Click **Add**.
- g) Repeat these steps to add additional attributes, as needed, for each authorization criteria.



Note: All criteria must pass in order for a user to be authorized.

- h) Click **Show Advanced Criteria** to configure more complex evaluations (for example, OR conditions) using OGNL expressions.



Note: Expressions must be enabled for the **Show Advanced Criteria** button to appear.

- Use the in-line editor to enter the OGNL expression.
- Use the **Error Result** field to enter an error message or an error code for use if authorization fails.



Note: If the OGNL expression resolves to a string value (rather than true or false), then the returned value overrides the error message defined in the **Error Result** field.

- Click **Add**. Continue to add expressions, as needed.
 - Click the **Test** link to input values and test the resulting output for the expression.
- i) Use the **Edit**, **Update**, and **Cancel** workflow to make or undo a change to an existing criterion. Use the **Delete** and **Undelete** workflow to remove an existing criterion or cancel the removal request.

5. On the **Summary** screen, click **Save** to complete the configuration.

Manage resource owner credential mappings

The **OAuth Settings > Resource Owner Credentials Mapping** workflow allows you to map attributes based on validated user credentials into the USER_KEY and extended attributes for a persistent grant. You may supplement credential-validation mapping sources with attribute look-ups from your user data source.

- To create a mapping, select the source of the attributes from the list and click **Add Mapping**.
- To modify an existing mapping, select it by its name under **Mappings**.

- To remove an existing mapping or to 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. In addition, any corresponding entries defined in the **OAuth Settings > Access Token Mapping** screen will also be removed.

Configure resource-owner attribute sources and user lookup

Attribute sources are specific data store or directory locations containing information that may be needed to fulfill the USER_KEY and any extended attribute values of persistent grants. This configuration is optional.

1. Click **Add Attribute Source** (or click **Next** if you are not using a data store).
2. On the **Data Store** screen, enter values into the **Attribute Source ID** and **Attribute Source Description** fields, select an data store from the list, and click **Next** to continue the configuration.

Define resource-owner contract fulfillment

On the **Contract Fulfillment** screen, you map values into the USER_KEY and extended attributes for a persistent grant. Use this mapping for the resource owner password credential grant type.



Note: The USER_KEY values must be unique across all end users because it is the identifier to store and to retrieve persistent grants. For example, if you have two Active Directory domains, the sAMAccountName value of an end user in one domain may collide with that of another end user in the other domain. In this scenario, you can map Subject DN to the USER_KEY.

Map each attribute from one of these Sources:

- Context

Values are returned from the context of the transaction at runtime.



Note: The HTTP Request 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. (If the **Expression** selection is not available, you may enable it by editing the `org.sourceid.common.ExpressionManager.xml` file in the `<pf_install>/pingfederate/server/default/data/config-store` directory.)

- Password Credential Validator

When you make this selection, the associated Value drop-down list consists of the attributes (for example, username) associated with the credential-validation instance.

- LDAP/JDBC/Custom (when a data store is used)

Values are returned from your data store (if used). When you make this selection, the Value list is populated by the LDAP, JDBC or Custom attributes identified for this data store.

- Expression (when enabled)

This option provides more complex mapping capabilities; for example, transforming incoming values into different formats. All of the variables available for text entries are also available for expressions.

- 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 syntax:

```
#{attribute}
```

You can also enter values from your data store, when applicable, using this syntax:

```
#{ds.attribute}
```

where `attribute` is any of the data store attributes you have selected.

1. Choose a source and then choose (or enter) a value for each attribute in the contract.
2. Click **Next**.

Specify issuance criteria for resource-owner credentials mapping


Use the **Issuance Criteria** screen to define criteria PingFederate® can evaluate to determine whether to issue an access token for a user. This token authorization can be used to restrict who can access an OAuth-protected resource.

To configure issuance criteria:

1. Select a source containing attributes for token authorization.

Associated attributes appear in the **Attribute Name** list:


- Context – Select to use values 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.

- LDAP/JDBC/Custom (if a data store is configured) – Select to access attributes returned from a data source.
- Mapped Attributes – Select to access the output of the mapping into the persistent grant—for example, the USER_KEY or an extended attribute.
- Password Credential Validator – Select to access attributes returned from the credential-validation instance.


2. Select an attribute name.

3. Select the condition you want to apply from the list.

 **Note:** Use one of the *multi-value* conditions when you want PingFederate to verify that an attribute contains or does not contain the specified value in its attribute list. Using a single-value condition when an attribute has multiple values causes the condition to evaluate consistently to false.

4. Enter the exact value in the **Value** field.

Wildcards (for partial matching) are not allowed.

 **Tip:** You can use OGNL expressions to perform more complex evaluations, including partial matching.


5. Optional: Use the **Error Result** field to customize an error code or error message for use if the token authorization fails.

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.


If you leave this field blank, the user sees a default ACCESS_DENIED error result at runtime if the authorization fails.

6. Click **Add**.


7. Repeat the steps above to add additional attributes, as needed, for each authorization criteria.

 **Note:** All criteria must pass in order for a user to be authorized.

8. Optional: Click **Show Advanced Criteria** to configure more complex evaluations (for example, OR conditions) using OGNL expressions.

 **Note:** Expressions must be enabled for the **Show Advanced Criteria** button to appear.

- Use the in-line editor to enter the OGNL expression.
- Use the **Error Result** field to enter an error message or an error code for use if authorization fails (see step 5 for more information).

 **Note:** If the OGNL expression resolves to a string value (rather than true or false), then the returned value overrides the Error Result.

- Click **Add**. Continue to add expressions, as needed.
- Click the **Test** link to input values and test the resulting output for the expression.

9. Click **Next** (or **Done**, if you are finished with the configuration, and then **Save** on the **Manage Mappings** screen).

To edit issuance criteria:

- Click **Edit** under Actions for the criteria.

To delete issuance criteria:

- Click **Delete** under Actions for the criteria and then click **Save**.

Review the resource-owner credentials mapping

When you have finished the configuration, you can review it on the **Summary** screen.

If you need to make any changes, click the heading over the information you want to edit. When you are finished, click **Done** and then **Save** on the **Manage Mappings** screen.


Manage access token mappings

In this required configuration, an administrator maps attributes to be requested from the OAuth resource server with the access token—the token attribute contract.


When mapping a Default context, you define how the OAuth AS maps values into the attributes based on the persistent-grant USER_KEY and any extended attributes.

When a specific context is selected, you can also map attributes from the selected context, namely the chosen IdP adapter, credential validator, or IdP connection (with an OAuth attribute mapping configuration) into the access tokens.

The mapping used at runtime depends on the authentication context of the original grant. If the authentication context results in a match, the OAuth AS uses that specific mapping; otherwise, it uses the default mapping for the applicable access token manager instance.

 **Note:** The **OAuth Settings > Access Token Mapping** workflow becomes available only after at least one instance of access token management has been configured.

- To create a mapping, select the source of the attributes from the **Context** list and the target from the **Access Token Manager** list, and then click **Add Mapping**.
- To modify an existing mapping, select it by its name under **Mappings**.
- To remove an existing mapping or to cancel the removal request, click **Delete** or **Undelete** under **Action**.

 **Note:** Before removing an existing mapping from your configuration, ensure that it is not used by your OAuth use cases.

Configure access token attribute sources and user lookup

Attribute sources are specific data store or directory locations containing information that may be needed to fulfill the token attribute contract. This configuration is optional.

1. Click **Add Attribute Source** (or click **Next** if you are not using a data store).
2. On the **Data Store** screen, enter values into the **Attribute Source ID** and **Attribute Source Description** fields, select an data store from the list, and click **Next** to continue the configuration.

Define access token attribute contract fulfillment

On the **Contract Fulfillment** screen, you map values into the token attribute contract. These are the attributes that will be included or referenced in the access token.

Map each attribute to fulfill the Token Attribute Contract from one of these Sources:

- Adapter/Password Credential Validator/IdP Connection

If you selected a specific IdP adapter, password credential validator or IdP connection under Context in the previous screen, you have the option to map attributes from that specific authentication system. Select the corresponding Context in this screen and choose the desired attribute under Value.

- Context

Values are returned from the context of the transaction at runtime.



Note: The HTTP Request 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. (If the **Expression** selection is not available, you may enable it by editing the `org.sourceid.common.ExpressionManager.xml` file in the `<pf_install>/pingfederate/server/default/data/config-store` directory.)

- Persistent Grant

When you make this selection, the associated Value drop-down list is populated by the `USER_KEY` and extended attributes from the persistent access-token grant.

- LDAP/JDBC/Custom (when a data store is used)

Values are returned from your user-data store. When you make this selection, the Value list is populated by the LDAP, JDBC or Custom attributes identified for this data store.

- Expression (when enabled)

This option provides more complex mapping capabilities; for example, transforming incoming values into different formats. All of the variables available for text entries are also available for expressions.

- 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 syntax:

```
#{USER_KEY}
```

You can also enter values from your data store, when applicable, using this syntax:

```
#{ds.attribute}
```

where `attribute` is any of the data store attributes you have selected.

1. Choose a source and then choose (or enter) a value for each attribute in the contract.
2. Click **Next**.

Specify issuance criteria for access token mapping

Use the **Issuance Criteria** screen to define criteria PingFederate® can evaluate to determine whether to issue an access token for a user. This token authorization can be used to restrict who can access an OAuth-protected resource. For example, before reissuing a token in a refresh-token scenario, the server can verify that the user is still current in your organization's user-data store and has retained the necessary permissions.

To configure issuance criteria:

1. Select a source containing attributes for token authorization.


Associated attributes appear in the **Attribute Name** list:

- Context – Select to use values 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.

- LDAP/JDBC/Custom (if a data store is configured) – Select to access attributes returned from a data source.
 - Mapped Attributes – Select to use access token attributes returned from token attribute contract mapping.
 - Persistent Grant – When you select this option, the associated **Attribute Name** list is populated by the `USER_KEY` and extended attributes from the persistent access-token grant.
2. Select an attribute name.
 3. Select the condition you want to apply from the list.

 **Note:** Use one of the *multi-value* conditions when you want PingFederate to verify that an attribute contains or does not contain the specified value in its attribute list. Using a single-value condition when an attribute has multiple values causes the condition to evaluate consistently to false.

4. Enter the exact value in the **Value** field.

Wildcards (for partial matching) are not allowed.


 **Tip:** You can use OGNL expressions to perform more complex evaluations, including partial matching.

5. Optional: Use the **Error Result** field to customize an error code or error message for use if the token authorization fails.


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.

If you leave this field blank, the user sees a default `ACCESS_DENIED` error result at runtime if the authorization fails.


6. Click **Add**.
7. Repeat the steps above to add additional attributes, as needed, for each authorization criteria.

 **Note:** All criteria must pass in order for a user to be authorized.

8. Optional: Click **Show Advanced Criteria** to configure more complex evaluations using OGNL expressions.

 **Note:** Expressions must be enabled for the **Show Advanced Criteria** button to appear.

- Use the in-line editor to enter the OGNL expression.
- Use the **Error Result** field to enter an error message or an error code for use if authorization fails (see step 5 for more information).

 **Note:** If the OGNL expression resolves to a string value (rather than true or false), then the returned value overrides the Error Result.

- Click **Add**. Continue to add expressions, as needed.
- Click the **Test** link to input values and test the resulting output for the expression.

9. Click **Next** (or **Done**, if you are finished with the configuration, and then **Save** on the **Manage Mappings** screen).

To edit issuance criteria:

- Click **Edit** under Actions for the criteria.

To delete issuance criteria:

- Click **Delete** under Actions for the criteria and then click **Save**.

Review the access token mapping

When you have finished the configuration, you can review it on the **Summary** screen.

If you need to make any changes, click the heading over the information you want to edit. When you are finished, click **Done** and then **Save** on the **Manage Mappings** screen.

Configure an OAuth assertion grant IdP connection

An OAuth assertion grant connection exchanges a SAML assertion or a JWT for an OAuth access token with the PingFederate® OAuth AS. You can configure an OAuth assertion grant connection with an IdP partner either in conjunction with browser-based SSO, WS-Trust, or independently.

For more information about these standards, see [Security Assertion Markup Language \(SAML\) 2.0 Profile for OAuth 2.0 Client Authentication and Authorization Grants](https://tools.ietf.org/html/rfc7522) (tools.ietf.org/html/rfc7522) and [JSON Web Token \(JWT\) Profile for OAuth 2.0 Client Authentication and Authorization Grants](https://tools.ietf.org/html/rfc7523) (tools.ietf.org/html/rfc7523).

1. On the **Connection Type** screen, select the **OAuth Assertion Grant** check box.



Tip: You may also select other options (for example, the **Browser SSO Profiles** check box). If you do, you will be prompted to complete the required configuration.

For simplicity, this topic only focuses on the **OAuth Assertion Grant** configuration.

2. On the **General Info** screen, enter the required information.
3. On the **OAuth Assertion Grant Attribute Mapping** screen, click **Configure OAuth Assertion Grant Attribute Mapping** to continue.

Define an attribute contract for the OAuth assertion grant

An attribute contract is a set of user attributes the IdP sends in the SAML assertion or JWTs for this connection. You identify these attributes on the **Attribute Contract** screen.

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.

To add an attribute:

1. Enter the attribute name in the text box.

Attribute names are case-sensitive and must correspond to the attribute names expected by your partner.

2. Optional: Select the check box under Mask Values in Log.
3. Click **Add**.

To modify an attribute name or masking selection:

1. Click **Edit** under Action for the attribute.
2. Make the change and click **Update**.



Note: If you change your mind, ensure that you click the Cancel link in the Actions column, not the Cancel button, which discards any other changes you might have made in the configuration steps.

To delete an attribute:

- Click **Delete** under Action for the attribute.

Configure access token manager mappings

Use the **Access Token Manager Mapping** screen to associate one or more access token manager instances with this connection to define how access tokens are created.

To create a new access token manager mapping:

- Click **Create New Access Token Manager Mapping**.

To edit an existing access token manager mapping:

- Click on the access token manager instance and click the step you need to change.


To delete an access token manager mapping:

1. Click **Delete** next to the access token manager instance. (To undo the deletion, click **Undelete**.)
2. Click **Save** to confirm the deletion.

Select an access token manager instance

Use the **Access Token Manager** screen to select an access token manager instance to associate with this connection. Attributes are mapped from the connection's contract into the access token's contract to define the resulting content of created access token.

To select an access token manager instance:

- Select an access token manager instance from the list.
 -  **Tip:** If the access token manager instance you need is not available, click **Manage Access Token Management Instances** to define one or more instances you need for this connection.

Configure a data store for OAuth assertion grant attribute mapping

This optional configuration is the same for all OAuth attribute-mapping task flows.

- If you do not need additional attributes from a data store, just click **Next** on the **Data Store** screen.

Define OAuth assertion grant contract fulfillment

On the **Contract Fulfillment** screen, map values from the SAML assertions or JWTs to the attributes defined for the attribute contract. These are the values that the Access Token Manager instance requires to create an OAuth access token.

At runtime, an SSO operation fails if PingFederate® cannot fulfill the required attribute.

Map attributes from one of the following Sources:


- **Assertion**

When selected, the **Value** list is populated 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 is populated with the available context of the transaction.

 **Note:** The **HTTP Request** context value is retrieved as a Java object rather than text. For this reason, OGNL expressions are more appropriate to evaluate and return values (see **Expression**).

- **LDAP, JDBC, or Custom**

When selected, the **Value** list is populated with attributes that you have selected from the data store. Select the desired attribute from the list.

- **Expression** (when enabled)

This option provides more complex mapping capabilities; for example, 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 (see **Text**).

Note that expressions are not enabled by default. For more information about enabling and editing OGNL expressions, see [Attribute mapping expressions](#) on page 497.

- **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 data store, when applicable, using this syntax:

```
${ds.attribute}
```

where *attribute* is any attribute that you have selected from the data store.

1. Choose a source and then choose (or enter) a value for each attribute in the contract.
2. Click **Next**.


Specify issuance criteria for OAuth assertion grant

Use the **Issuance Criteria** screen to define criteria PingFederate® can evaluate to determine whether to issue an access token for a user. This token authorization can be used to restrict who can access an OAuth-protected resource.

To configure issuance criteria:


1. Select a source containing attributes for token authorization.

Associated attributes appear in the **Attribute Name** list:

- Assertion – Select to access attributes from the assertion.
- Context – Select to use values 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.
- LDAP/JDBC/Custom (if a data store is configured) – Select to access attributes returned from a data source.
- Mapped Attributes – Select to access the output of the mapping into the access token.


2. Select an attribute name.

3. Select the condition you want to apply from the list.

 **Note:** Use one of the *multi-value* conditions when you want PingFederate to verify that an attribute contains or does not contain the specified value in its attribute list. Using a single-value condition when an attribute has multiple values causes the condition to evaluate consistently to false.

4. Enter the exact value in the **Value** field.

Wildcards (for partial matching) are not allowed.

 **Tip:** You can use OGNL expressions to perform more complex evaluations, including partial matching.


5. Optional: Use the **Error Result** field to customize an error code or error message for use if the token authorization fails.

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.


If you leave this field blank, the user sees a default ACCESS_DENIED error result at runtime if the authorization fails.

6. Click **Add**.


7. Repeat the steps above to add additional attributes, as needed, for each authorization criteria.

 **Note:** All criteria must pass in order for a user to be authorized.

8. Optional: Click **Show Advanced Criteria** to configure more complex evaluations using OGNL expressions.

 **Note:** Expressions must be enabled for the **Show Advanced Criteria** button to appear.

- Use the in-line editor to enter the OGNL expression.
- Use the **Error Result** field to enter an error message or an error code for use if authorization fails (see step 5 for more information).

 **Note:** If the OGNL expression resolves to a string value (rather than true or false), then the returned value overrides the Error Result.

- Click **Add**. Continue to add expressions, as needed.
- Click the **Test** link to input values and test the resulting output for the expression.

9. Click **Next** (or **Done**, if you are finished with the configuration, and then **Save** on the **Manage Mappings** screen).

To edit issuance criteria:

- Click **Edit** under Actions for the criteria.

To delete issuance criteria:

- Click **Delete** under Actions for the criteria and then click **Save**.

Review OAuth assertion grant attribute mapping configuration

When you finish the configuration, you can review it on the **Summary** screen.

If you need to make any changes, click the heading over the information you want to edit. When you finish, click **Done**.

Review OAuth assertion grant configuration

When you finish the configuration, you can review it on the **Summary** screen.

If you need to make any changes, click the heading over the information you want to edit. When you finish, click **Done**.

Configure OpenID Connect policies

This configuration allows you to define OpenID Connect policies for client access to attributes mapped according to OpenID specifications. It also provides an option to include a session identifier in the ID Tokens, which could be useful for the relying parties, such as PingAccess® for client session management.

- To configure a new policy, click **Add Policy**.

On the **Manage Policy** screen, enter the required information and configure optional settings as follows:


1. Enter a policy ID and name.

Note that the policy ID cannot be changed once the policy is saved.

2. Select an access token manager instance from the list.
3. Change the default value of the ID token lifetime in minutes.

The default value is 5 and can range from 0 to 65,535.


4. Select the **Include Session Identifier in ID Token** check box to add a session identifier in the ID tokens.
5. Select the **Include User Info in ID Token** check box to include additional attributes in the ID tokens.

 **Tip:** Alternatively, OAuth clients can obtain additional attributes from the UserInfo endpoint (`/idp/userinfo.openid`).

- To modify an existing policy, select it by its name under **Policy ID** and make changes on the **Manage Policy** screen.
- To review the usage of an existing policy, click **Check Usage** under **Action**.
- To remove an existing policy or to cancel the removal request, click **Delete** or **Undelete** under **Action**.
- To elect an existing policy to be the default OpenID Connect policy, click **Set as Default** under **Action**.

Configure the policy attribute contract

On the **Attribute Contract** screen, specify attributes you want returned from the PingFederate® UserInfo endpoint for OpenID Connect (`/idp/userinfo.openid`).

 **Note:** The `sub` value of this configuration is also used to fulfill an internally managed ID Token attribute contract. The contract usage is dynamic, depending on how and when the user authenticated.

To add an attribute:

- Enter the attribute name in the text box and click **Add**.

To modify an attribute name:

1. Click **Edit** under Action for the attribute.
2. Make the change and click **Update**.

To delete an attribute:

- Click **Delete** under Action for the attribute.

Configure policy attribute sources and user lookup

Attribute sources are specific data store or directory locations containing information that may be needed to fulfill the attribute contract for a policy. This configuration is optional.

1. Click **Add Attribute Source** (or click **Next** if you are not using a data store).
2. On the **Data Store** screen, enter values into the **Attribute Source ID** and **Attribute Source Description** fields, select an data store from the list, and click **Next** to continue the configuration.

Define policy contract fulfillment

On the **Contract Fulfillment** screen map attributes from the access token or other sources to fulfill the attribute contract.

Map the subject attribute and all extended attributes from one of these Sources:

- Context

Values are returned from the context of the transaction at runtime.



Note: The HTTP Request 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. (If the **Expression** selection is not available, you may enable it by editing the `org.sourceid.common.ExpressionManager.xml` file in the `<pf_install>/pingfederate/server/default/data/config-store` directory.)

- LDAP/JDBC/Custom (when a data store is used)

Values are returned from your data store (if used). When you make this selection, the Value list is populated by the LDAP, JDBC or Custom attributes identified for this data store.

- Expression (when enabled)

This option provides more complex mapping capabilities; for example, transforming incoming values into different formats. All of the variables available for text entries are also available for expressions.

- 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 syntax `${attribute}`.

You can also enter values from your data store, when applicable, using this syntax:

```
${ds.attribute}
```

where `attribute` is any of the data store attributes you have selected.

- Access Token

The value is provided from the access token.

1. Choose a source and then choose (or enter) a value for each attribute in the contract.
2. Click **Next**.


Specify issuance criteria for policy mapping

Use the **Issuance Criteria** optional screen to define criteria PingFederate® can evaluate to determine whether to issue the ID token and access token. If a criterion fails, no tokens are issued and access is denied.

To configure issuance criteria:


1. Select a source containing attributes for token authorization.

Associated attributes appear in the **Attribute Name** list:

- Context – Select to use values 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.
- LDAP/JDBC/Custom (if a data store is configured) – Select to access attributes returned from a data source.
- Mapped Attributes – Select to access the output data returned from UserInfo endpoint.
- Access Token – Select to use attributes from the access token.


2. Select an attribute name.

3. Select the condition you want to apply from the list.

-  **Note:** Use one of the *multi-value* conditions when you want PingFederate to verify that an attribute contains or does not contain the specified value in its attribute list. Using a single-value condition when an attribute has multiple values causes the condition to evaluate consistently to false.

4. Enter the exact value in the **Value** field.

Wildcards (for partial matching) are not allowed.

-  **Tip:** You can use OGNL expressions to perform more complex evaluations, including partial matching.


5. Optional: Use the **Error Result** field to customize an error code or error message for use if the token authorization fails.

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.


If you leave this field blank, the user sees a default ACCESS_DENIED error result at runtime if the authorization fails.

6. Click **Add**.


7. Repeat the steps above to add additional attributes, as needed, for each authorization criteria.

-  **Note:** All criteria must pass in order for a user to be authorized.

8. Optional: Click **Show Advanced Criteria** to configure more complex evaluations (for example, OR conditions) using OGNL expressions.

-  **Note:** Expressions must be enabled for the **Show Advanced Criteria** button to appear.

- Use the in-line editor to enter the OGNL expression.
- Use the **Error Result** field to enter an error message or an error code for use if authorization fails (see step 5 for more information).

-  **Note:** If the OGNL expression resolves to a string value (rather than true or false), then the returned value overrides the Error Result.

- Click **Add**. Continue to add expressions, as needed.
- Click the **Test** link to input values and test the resulting output for the expression.

9. Click **Next** (or **Done**, if you are finished with the configuration, and then **Save** on the Manage Policies screen).

To edit issuance criteria:

- Click **Edit** under Actions for the criteria.

To delete issuance criteria:

- Click **Delete** under Actions for the criteria and then click **Save**.

Review the policy

When you have finished the configuration, you can review it on the **Summary** screen.

If you need to make any changes, click the heading over the information you want to edit. When you are finished, click **Done** and then **Save** on the **Manage Policies** screen.



Note: If this is the first policy you are creating, you must designate it as the default before saving. You can change the default as needed when you create additional policies.

Manage OAuth clients

An OAuth client application interacts with the PingFederate® OAuth AS to obtain access tokens needed to call OAuth-protected services at the Resource Server (RS).



Important: OAuth client records, like other PingFederate configuration data, are maintained in XML files. Alternatively, you can define an OAuth client data store (a centralized database) to store client records.

The **OAuth Settings** menu displays a list of the most-recently modified OAuth clients. As needed, you can access the rest of your clients on the **OAuth Settings > Client Management** screen.

- To add a client, click **Add Client** and complete the configuration on the **Clients** screen.
- To edit a recently modified client, select it by its name and update its configuration on the **Clients** screen.
- To access the rest of your clients, click **Manage All** to open the **Client Management** screen.

The **Client Management** screen displays 20 OAuth clients at a time. Use the pagination controls to navigate through the rest of your clients. You can also search clients by their names or client IDs.



Tip: A client is included in the search results so long as its name *or* ID is a partial, case-insensitive match to a search term.

On the **Client Management** screen, you can sort by client name or ID.

- To remove a client or cancel the removal request, use the **Delete** and **Undelete** workflow next to the applicable client, and then click **Save**.




Configure an OAuth client



The **Client** screen provides controls over the usage and behavior of the applications requesting access to protected resources through the PingFederate® OAuth AS.



Configure the OAuth client to suit your use cases.

Refer to the following table for detailed information about each field.

Field	Description
Client ID (Required)	A unique identifier the client provides to the Resource Server (RS) to identify itself. This identifier is included with every request the client makes.
Name (Required)	A descriptive name for the client instance. This name appears when the 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.

Field	Description
	<ul style="list-style-type: none"> • Select None 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: <ul style="list-style-type: none"> • This client uses the Client Credentials grant type (see the Allowed Grant Types field). • 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 endpoint (see the Grant Access to Session Revocation API field). • Select Client Secret for HTTP Basic authentication. <ul style="list-style-type: none"> • Click Generate Secret to create a strong random alphanumeric string or manually enter a 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. • Select Client TLS Certificate for mutual SSL/TLS authentication; recommended for client applications where security policies prohibit storing passwords. <ul style="list-style-type: none"> • Select a Trusted CA from the Issuer list. (These are CA certificates imported into PingFederate.) Alternatively, you may choose to trust any of the issuers listed 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.  Note: Mutual SSL/TLS authentication requires the use of a secondary PingFederate SSL port. • Select Private Key JWT check box if the client authenticates via the <code>private_key_jwt</code> client authentication method, as defined in <i>Client Authentication</i> in the OpenID Connect specification (openid.net/specs/openid-connect-core-1_0.html#ClientAuthentication). <ul style="list-style-type: none"> • Select the Replay Prevention check box if your case requires a unique JWT for each request. When selected (not by default), the underlying Assertion Replay Prevention Service handles the validation of the uniqueness requirement.  Note: The Assertion Replay Prevention Service is cluster-aware. For further information, see <i>Assertion Replay Prevention Service</i> on page 546.
Require Signed Request	Select this check box if the client submits request parameters as claims in a self-contained, signed JWT request object, as defined in the OpenID Connect specification (see the section about passing a request object by value at http://openid.net/specs/openid-connect-core-1_0.html#RequestObject).
JWKS URL, and JWKS	<p>If the client uses signed JWTs for client authentication or transmission of request parameters (or for both purposes), the partner must provide the corresponding JSON Web Key (JWK) or a set of JWKS so that PingFederate can validate the integrity of the JWTs.</p> <p>If the partner provides an URL, enter it in the JWKS URL field; otherwise, enter the actual JWKS in the JWKS text field.</p>

Field	Description
Redirection URIs	<p>URIs to which the OAuth AS may redirect the resource owner's user agent after authorization is obtained. A redirection URI is used with the Authorization Code and Implicit grant types.</p> <p>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:</p> <p><code>https://www.example.com/OAuthClientApp/callback.jsp</code></p> <p> 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 <code>redirect_uri</code> parameter when contacting the authorization endpoint (<code>/as/authorization.oauth2</code>) and token endpoint (<code>/as/token.oauth2</code>).</p>
Logo URL	<p>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.)</p>
Bypass Authorization Approval	<p>When selected, resource-owner approval for client access is assumed—the authorization consent page is not presented to the user for this client.</p> <p>Use this setting, for example, when you want to deploy a trusted application and authenticate end users via an IdP adapter or IdP connection.</p>
Restrict Scopes	<p>Restricts this client's access to specific scopes or scope groups.</p> <p>When selected, a list appears of all scopes defined within the PingFederate AS. Select the scopes or scope groups available for this client.</p> <p> Note: Selecting this box and not selecting any specific scopes restricts this client's access to only the default scope defined in the OAuth Settings > Authorization Server Settings screen. If a client requests a specific scope (not selected in the list), an error is returned.</p>
Allowed Grant Types	<p>Select at least one grant type that this client is allowed to use. (There is no default selection.)</p> <p>Available grant types are:</p> <ul style="list-style-type: none"> • Authorization Code • Resource Owner Password Credentials • Refresh Token • Implicit • Client Credentials • Access Token Validation (Client is a Resource Server) • Extension Grants <p>(For more information about each grant type, see Grant types.)</p>
Default Access Token Manager	<p>Select a default access token management (ATM) instance for this client.</p>
Validate Against All Eligible Access Token Managers	<p>Applicable only to resource server (RS) clients.</p> <p>If selected, this RS client is not required to specify additional parameters (<code>access_token_manager_id</code> or <code>aud</code>) to disambiguate the ATM instance in its token validation requests. When the RS client does not specify the desired ATM instance, PingFederate validates the access tokens against all eligible ATM instances. This</p>

Field	Description
	<p>simplifies interactions with PingAccess® by avoiding the need to align resource URIs between PingAccess and PingFederate.</p> <p>This check box is not selected by default.</p>
Persistent Grants Expiration	<p>Allows an administrator to override the Persistent Grant Lifetime field value set globally in the OAuth Settings > Authorization Server Settings screen.</p> <p>Options are:</p> <ul style="list-style-type: none"> • Use Global Setting (the default selection) • Grants Do Not Expire • Customer value in days, hours, or minutes.
Refresh Token Rolling Policy	<p>Allows an administrator to override the Roll Refresh Token Values setting configured globally in the OAuth Settings > Authorization Server Settings screen.</p> <p>Options are:</p> <ul style="list-style-type: none"> • Use Global Setting (the default selection) • Roll <p>Note that this selection does not override the Minimum Interval to Roll Refresh Tokens (Hours) value set on the Authorization Server Settings screen.</p> <ul style="list-style-type: none"> • Don't Roll
OpenID Connect	<p> Note: These options are displayed only when the OpenID Connect role is enabled in Server Configuration > Server Settings > Roles & Protocols screen.</p> <p>ID Token Signing Algorithm: Select the signing algorithm for the ID tokens. The default algorithm is RSA using SHA-256.</p> <p>Policy: Choose a specific OpenID Connect policy.</p> <p>Grant Access to Session Revocation API: Select this check box to allow this client application to add sessions to or query the revocation status via the Back-Channel Session Revocation API endpoint at <code>/pf-ws/rest/sessionMgmt/revokedSris</code>. Authentication is required. This check box is not selected by default.</p> <p> Note: If the Track User Sessions for Logout check box is selected in the OAuth Settings > Authorization Server Settings screen, there are two additional fields: PingAccess Logout Capable and Logout URIs.</p> <p>PingAccess Logout Capable: When selected, PingFederate sends (via the browser) logout requests to an OpenID Connect endpoint in PingAccess as part of the logout process (see OpenID Connect endpoints in the PingAccess documentation). This check box is not selected by default.</p> <p>Logout URIs: Enter additional endpoints at the relying parties as needed. PingFederate sends (via the browser) requests to these URIs as part of the logout process.</p>

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 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.

 **Tip:** 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, PingFederate® can be configured to send (via the browser) logout requests to PingAccess® and additional requests to other relying parties.

When the PingAccess option is selected, PingFederate sends logout requests (via 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 SP-initiated logout request reaches the PingFederate IdP server, the same logout process applies as well. Depending on the enterprise architecture, this could further improve single sign-on and logout use cases.

Configure Asynchronous Front-Channel Logout

1. Configure Asynchronous Front-Channel Logout settings in the **OAuth Settings > Authorization Server Settings** screen.
 - a) Select the **Track User Sessions for Logout** check box to enable Asynchronous Front-Channel Logout for OpenID Connect clients
 - b) Optional: Select the **Revoke User Session on Logout** check box to add the PingFederate session to the revocation list.
2. Optional: Configure logout settings for each applicable OAuth client.
 - a) Select the **PingAccess Logout Capable** check box to send (via the browser) logout requests to PingAccess.
 - b) Enter additional logout endpoints at the relying parties under **Logout URIs** to send logout requests to close sessions on their end.

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.


Configure Back-Channel Session Revocation

1. For each applicable OAuth client, select the **Grant Access to Session Revocation API** check box in its client configuration screen.
2. For each OpenID Connect policy that has been applied to OAuth clients supporting session revocation, select the **Include Session Identifier in ID Token** check box.

OpenID Connect policies are defined in the **OAuth Settings > OpenID Connect Policy Management** screen.

Alternatively, if authentication sessions are enabled, you may enable session validation for the applicable access token management instance such that the session identifier becomes available through the access tokens. OpenID Connect is not required in this case. For more information, see [Manage session validation settings](#) on page 250.

3. Optional: On the **OAuth Settings > Authorization Server Settings** screen, modify the **Session Revocation Lifetime** value.

 **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 **Session Revocation Lifetime** value of 250 minutes exceeds the default lifetime of the PingAccess (PA) Tokens by 10 minutes to allow for clock skew between servers (see [Web Sessions](#) in the PingAccess documentation).

OAuth attribute mapping using a data store

This optional configuration is the same for all of the OAuth attribute-mapping task flows, including:

- [Configure OpenID Connect policies](#) on page 266
- [Manage resource owner credential mappings](#) on page 257
- [Manage IdP adapter mappings for OAuth](#) on page 253
- [Manage access token mappings](#) on page 260


A similar configuration is also used for attribute mapping in an IdP connection (see [Configure OAuth attribute mapping](#) on page 354) and when configuring an OAuth Assertion Grant connection (see [Configure an OAuth assertion grant IdP connection](#) on page 262).

(For additional information, see [Fulfillment by data store queries](#) on page 506.)

Choose a data store

On the **Data Store** screen, choose a data store for PingFederate® to look up attributes.

1. Enter a description (and ID if prompted) for the data store.
2. Select a data store instance from the **Active Data Store** list.

 **Tip:** 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.

3. Depending on the data store type, the rest of the setup varies as follows:

Data store type	Required tasks
JDBC	<ul style="list-style-type: none"> • Specify a JDBC database table and columns on page 508 • Enter a database search filter on page 509
LDAP	<ul style="list-style-type: none"> • Specify an LDAP directory and attributes on page 510 • Define encoding for LDAP binary attributes on page 511 (optional) • Enter an LDAP search filter on page 511

Data store type	Required tasks
Custom	<ul style="list-style-type: none"> • Specify a custom source filter and fields on page 512

Identity provider SSO configuration

In an IdP role, you use the PingFederate® administrative console to configure local application-integration information and to manage connections to your SP-partner sites. Prior to configuring connections to SPs, you must establish your site as an IdP on the **Server Configuration > Server Settings > Roles & Protocols** screen.

Note that only one connection is needed 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 **Server Configuration > Server Settings > Federation Info** screen, you may identify your organization differently through the use of virtual server IDs on a per-connection basis (see [Multiple virtual server IDs](#) on page 84).

Additionally, you may deploy an SP connection to bridge a service provider to one or more identity providers through one or more authentication policy contracts (see [Federation hub use cases](#) on page 86).



Note: This topic applies to configuration settings needed for browser-based SSO. While there is some cross-over information also applicable to WS-Trust STS, if you are using PingFederate *exclusively* as an STS, start with [WS-Trust STS configuration](#) on page 400.

IdP application integration settings

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 Software Development Kit (see [SSO integration kits and adapters](#) on page 67).

Under **Application Integration** on the **IdP Configuration** menu, you configure instances of the IdP adapters (for example, the HTML Form Adapter) that PingFederate needs to interact with applications or access-management systems used to authenticate users at your site. You can also set a Default URL to which users may be directed during SLO, and you can 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 STS, the selections under **Application Integration** also include menu items for configuring plug-in **Token Processors** and optionally **STS Request Parameters** (see [IdP configuration for STS](#)).

Manage IdP adapters

An IdP adapter is used to look up session information and provide user identification to PingFederate®. You must configure at least one instance of an IdP adapter in order to set up connections to SP partners.

PingFederate comes bundled with the following adapters:

- Kerberos Adapter
- HTML Form Adapter
- HTTP Basic Adapter
- OpenToken Adapter
- Composite Adapter

You can also deploy additional integration kits from the Ping Identity [Downloads](#) website.

You manage IdP adapter instances on the **IdP Configuration > Adapters** screen.

- 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**.
- ❗ **Important:** After installing new adapter program files, you may be required to make additional configuration changes in areas such as adapter instances and connections, as prompted by the administrative console.
- 📄 **Note:** Automatic multi-connection error checking occurs by default whenever you access this screen. The intent is to verify that configured connections have not been adversely affected by changes made here.
If you experience noticeable delays in accessing this screen, you can optionally disable automatic connection validation on the **Server Configuration > Server Settings > System Options** screen.

Create an IdP adapter instance

The first step in creating an adapter instance is choosing an adapter type.

- On the **Type** screen, configure the basics of this adapter instance.
 - a) Enter the required information and select the adapter type from the 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 screens.

Configure an IdP adapter instance

Depending on the selected adapter, the **IdP Adapter** screens present you with different configuration parameters.

- Follow the on-screen instructions to configure the adapter instance.
 - 📄 **Note:** If this is a child instance, select the override check box to modify the configuration.
- If you have deployed an adapter from an integration kit, locate the user guide from our [documentation website](#) and configure the adapter instance accordingly.
- Click **Show Advanced Fields** to enter an authentication context URI for an adapter that supports this feature.
 - 📘 **Tip:** Standard URIs are defined in the SAML specifications (see the OASIS documents [oasis-sstc-saml-core-1.1.pdf](#) and [saml-authn-context-2.0-os.pdf](#)).

Applicable only if the selected adapter supports the notion of authentication context.

Invoke IdP adapter actions

Adapters may be written to provide configuration assistance or validation actions. Actions may also include generation of parameters that might need to be set manually in a configuration file.

- Follow the on-screen instructions to complete the actions required.

Extend an IdP adapter contract

You can extend the IdP adapter contract with additional attributes in the **Extended Contract** screen. For the Composite Adapter, attributes from the IdP adapter instances that comprise the composite configuration *must* be added on this screen.

If the adapter does not return values for the extended attributes (or if you prefer to fulfill them differently using data store queries, dynamic text values, or results from OGNL expressions), you can define their fulfillment using the **Adapter Contract Mapping** workflow later.


- 📄 **Note:** If this is a child instance, select the override check box to modify the configuration.

- Enter the name of the desired attribute and click **Add**.

Repeat this step as needed to add another attribute.


Set pseudonym and masking options

- On the **Adapter Attributes** screen, configure the pseudonym and masking options.

 **Note:** The **Override Attributes** check box in this screen reflects the status of the override option in the **Extended Contract** screen.

- 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 SP partners use pseudonyms for account linking.


 **Note:** A selection is required regardless of whether 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 (for example, email) to prevent the same pseudonym from being assigned 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

Define the IdP adapter contract

An IdP adapter contract is a contract that may be used to fulfill the attribute contract passed to your SP partners. By default, PingFederate® fulfills the IdP adapter contract with attribute values from the adapter. Optionally, you can configure PingFederate to fulfill the IdP adapter contract with attribute values from local data stores, dynamic text values, results from OGNL expressions, or a combination of them. In addition, you have the option to verify requests using the token authorization workflow.

To customize the IdP adapter contract,

1. Select the applicable IdP adapter instance.
2. Go to the **Adapter Contract Mapping** screen.
3.  **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** screen, in which case the **Override Adapter Contract** check box is automatically selected for you.

Click **Configure Adapter Contract**.

Define attribute sources and user lookup

Attribute sources are specific data store or directory locations containing information that may be needed for the IdP adapter contract or the token authorization workflow. You can use more than one attribute source when mapping values to the IdP adapter contract.

The PingFederate IdP server supports separate data stores to look up attributes for a single mapping. For example, you can query multiple data stores for values and map those values in one mapping, or query a data store for a value and use that value as input for subsequent queries of other data stores.

Queries are executed in the order as shown on the **Attribute Sources & User Lookup** screen. Use the up and down arrows as needed to adjust the order.

If a required attribute (such as the user identifier `username` or `subject` for the HTML Form Adapter or the OpenToken IdP Adapter, respectively) cannot be fulfilled, the request fails.

- If your use case requires only dynamic texts or results from OGNL expressions without any attributes from local data stores, skip to the next screen.
- To add an attribute source, click **Add Attribute Source**.

Repeat this step as needed to add another attribute source.



- To modify an existing instance, select it by its name under **Description**.
- To remove an existing instance or to cancel the removal request, click **Delete** or **Undelete** under **Action**.

Fulfill an adapter contract

From the **Adapter Contract Fulfillment** screen, map values into the IdP adapter contract.

1. For a given attribute, select a source from the list.

For more information about the **Source** list, refer to the following table:

Source	Description
Adapter	Values are returned from the IdP adapter without any customization.
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.
JDBC, LDAP, or Custom (if configured)	Values are returned from your attribute source. When you make this selection, the Value list is populated by the JDBC, LDAP, or Custom attributes you identified for this Attribute Source.
Expression (when enabled)	This option provides more complex mapping capabilities; for example, transforming incoming values into different formats. All of the variables available for text entries are also available for expressions.  Tip: If you have not already done so, you may enable OGNL expression by editing the <code>org.sourceid.common.ExpressionManager.xml</code> file in the <code><pf_install>/pingfederate/server/default/data/config-store</code> directory. Restart PingFederate after saving the change. For a clustered PingFederate environment, edit the <code>org.sourceid.common.ExpressionManager.xml</code> file on the console node, sign on to the administrative console to replicate this change to all engine nodes on the Server Configuration > Cluster Management screen, and restart all nodes.
Text	The value is what you enter. You can enter text or mix text with references to any of the values from the IdP adapter, using the <code>\${attribute}</code> syntax. You can also enter values from your data store, when applicable, using this syntax: <code>\${ds.attr-source-id.attribute}</code> where <code>attr-source-id</code> is the Attribute Source Id value in the Data Store screen and <code>attribute</code> is any of the data store attributes you select. 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. Specify a value.

Not applicable when **Adapter** is the chosen from the **Source** list.

3. Repeat these steps until all attributes in the IdP adapter contract are mapped.

Define issuance criteria for adapter contract

Use the **Issuance Criteria** screen to define criteria that PingFederate® can evaluate to determine whether to approve or reject requests for protected resource.


The configuration begins by selecting a source containing the attribute to be verified. Some sources, such as **Context** and **Mapped Attributes**, are common to almost all use cases. Other sources, such as **Assertion**, have dependency on the type of configuration, for which the administrative console automatically hides the irrelevant sources automatically. After selecting a source, choose an attribute to be verified. Depending on the selected source, attributes vary. For example, if **Account Linking** is the selected source, **Local User ID** is the available attribute.

Finally, select the comparison method, specify the desired value, and enter an optional error message for scenarios where the condition fails.


When multiple criteria are configured through the user interface, all criteria must be met at runtime (evaluated as *true*) for the request to succeed; otherwise PingFederate rejects the request and returns an error message. If only a subset or a layer of the criteria is required, use one or more OGNL expressions to define them.

1. Select a source containing the attribute to be verified.


For more information about the **Source** list, refer to the following table:

Source	Description
Adapter	Select to access attributes from the IdP Adapter.
Context	Select to use values 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.
JDBC, LDAP, or Custom (if configured)	Select to access attributes returned from a data source.
Mapped Attributes	Select to access the mapped attributes.

2. Select an attribute from the **Attribute Name** list.
3. Select a comparison method from the **Condition** list.

 **Note:** Use multi-value conditions when you want PingFederate to verify that an attribute contains or does not contain the specified value in its attribute list. Using a single-value condition when an attribute has multiple values causes the condition to evaluate consistently to false.

4. Enter an exact value in the text field.

 **Note:** If your use case requires more complex evaluations, such as partial matching, use one or more OGNL expressions to define them.

5. Optional: Enter a custom error message in the **Error Result** field that PingFederate returns when the criterion fails.

If you leave this field blank, PingFederate returns `ACCESS_DENIED` at runtime when the criterion fails.


6. Click **Add**.


Use the **Edit**, **Update**, **Cancel**, **Delete**, and **Undelete** links to manage the criterion.

7. Optional: Repeat these steps to specify additional criteria.

Define issuance criteria using OGNL expressions

1. Click **Show Advanced Criteria**.

 **Note:** Expressions must be enabled for the **Show Advanced Criteria** button to appear.

 **Tip:** If you have not already done so, you may enable OGNL expression by editing the `org.sourceid.common.ExpressionManager.xml` file in the `<pf_install>/pingfederate/server/default/data/config-store` directory. 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 on the **Server Configuration > Cluster Management** screen, and restart all nodes.

2. Specify a criterion.

a) Enter an OGNL expression into the text field.

b) Enter a custom error message in the **Error Result** field.



Note: If the OGNL expression resolves to a string value (rather than *true* or *false*), the returned value overrides the custom error message.

c) Click **Add**.

d) Click the **Test** link to input values and test the resulting output for the expression.

Use the **Edit**, **Update**, **Cancel**, **Delete**, and **Undelete** links to manage the OGNL expression.

3. Optional: Repeat these steps to specify another OGNL expression.

If multiple OGNL expressions are specified, they all must pass in order for the request to succeed. Alternatively, you can combine multiple criteria in one OGNL expression.

Review an adapter contract

- On the **Summary** screen, review your configuration, modify as needed, and click **Done** to exit the **Adapter Contract Mapping** workflow.

Review and save an IdP adapter configuration

1. On the **Summary** screen, review your configuration, modify as needed, and click **Done** to exit the **Create Adapter Instance** workflow.
2. On the **Manage IdP Adapter Instances** screen, click **Save** to retain the configuration of the adapter instance.

If you want to exit without saving the configuration, click **Cancel**.

Configure a default URL and error message

As an IdP, you can specify to prompt end users to confirm their single logout requests and a default URL indicating a successful SLO to the end-user (if no other page is designated). On the **IdP Default URL** screen, 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 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.



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 [Localize messages for end users](#) on page 139. If localization is not needed, you may also specify a message directly in this field to change the default.

Your application or your partner's application may supply the URL at runtime (see [IdP endpoints](#) on page 448). However, if none is provided, PingFederate will use the default value you enter on this screen.



Tip: If you leave the default URL blank, PingFederate provides built-in landing page for the user. This web page is among the templates you can modify with your own branding or other information (see [Customizable user-facing screens](#) on page 126).

View IdP application endpoints

Web-application developers at your site need to know the application endpoints to initiate transactions via PingFederate.

- Click **Application Endpoints** on the **IdP Configuration** menu 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 448 and [System-services endpoints](#) on page 460.

View IdP protocol endpoints


Click **Protocol Endpoints** on the **IdP Configuration** menu to see a list of applicable SAML, WS-Federation, and WS-Trust STS endpoints. The pop-up window displays only those endpoints related to the federation protocols enabled on the **Server Configuration > Server Settings > Federation Info** screen. These endpoints are built into PingFederate® and cannot be changed.


Your federation partners or 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 favorite icon for all protocol endpoints. For more information, see [Customize the favicon for application and protocol endpoints](#) on page 144.

The table below 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.ssaml2 The SOAP endpoint that processes artifacts returned from a federation partner to retrieve the referenced XML message on the back channel. (See Important 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 Important note at the end of this table.)
Metadata Service	/ The default endpoint (empty path) from which partners can retrieve Auto-Connect metadata .
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 SSO requests.
Artifact Resolution Service (SAML 1.x)	/idp/soap.ssaml1 The SOAP endpoint that processes artifacts returned from a federation partner to retrieve the referenced XML message on the back channel. (See Important note at the end of this table.)
Single Sign-on Service (WS-Federation)	/idp/prp.wsf The WS-Federation implementation URL that receives and processes security-token requests and SLO messages.
WS-Trust STS (two endpoints)	/idp/sts.wst The SOAP endpoint that receives and processes security-token requests from STS clients (Web Service Clients at the IdP site) to be exchanged for a SAML token based on the configured SP connection. /pf/sts.wst

Service	URL and Description
	<p>Initiates direct STS token-to-token exchange and token validation from an IdP token processor to an SP token generator, when that feature is configured (see Token translator mappings on page 423).</p> <p> Note: If multiple token-processor instances of the same type are configured for the same connection or token-to-token mapping, a query parameter, <code>TokenProcessorId</code>, must be added to either of these endpoints—see Manage token processors on page 401.</p> <p>(See also “Important” footnote in this table.)</p>

 **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` (see [Configure PingFederate properties](#) on page 93).

Virtual server ID support

For SAML connections using multiple virtual server IDs (see [Multiple virtual server IDs](#) on page 84), each virtual server ID has its own set of protocol endpoints. You may export a connection metadata for your partner on the **Server Configuration > Metadata Export** screen (see [Provide SAML metadata by file](#) on page 116).

For WS-Federation (and SAML) connections using multiple virtual server IDs, you may provide your partner the federation metadata endpoint (`/pf/federation_metadata.ping`) with the `PartnerSpId` and `vsid` parameters; for example:

Partner's entity ID	Your virtual server ID	Federation metadata URL
SP	idev1	<code>https://www.example.com/pf/federation_metadata.ping?PartnerSpId=SP&vsid=idev1</code>
	idev2	<code>https://www.example.com/pf/federation_metadata.ping?PartnerSpId=SP&vsid=idev2</code>

(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 may provide your partner the STS metadata endpoint (`/pf/sts_mex.ping`) with the `PartnerSpId` 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 460.)

Note that the virtual server ID concept does not apply to the `/pf/sts.wst` endpoint because token-to-token exchange does not involve any connections. As needed, you may pass the token-to-token endpoint to your partners as-is.

Manage SP connections

As an IdP, you manage connection settings to support the exchange of federation-protocol messages (SAML, WS-Federation, or WS-Trust) with an SP or STS client application at your site.

These settings include:

- User attributes that you expect to send in an SSO token (SAML assertion, WS-Trust STS SAML token, or WS-Federation JSON Web Token).
- User attributes that may be sent using the SAML Attribute Query profile (if that profile is used).

- The protocol, profiles, and bindings of the connection, including detailed security specifications (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 (see [Federation planning checklist](#) on page 82).

If your agreement includes sending assertions containing attribute values from local data stores, you must define the required data stores (see [Manage data stores](#) on page 158).

Administrative interface

You manage connection settings using the **SP Connection** wizard, 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 screen, where you manage one or more settings. You may move to a sibling task using the **Next** or **Previous** button. If you are on a sub task, you may also move to its parent task using the **Done** button.

When creating a new connection, you may save your progress using the **Save Draft** button. Note that not all screens offer this option. When you reach the **Activation & Summary** screen, you must click **Save** to complete the new connection.

When editing an existing connection, you may make changes and then click **Save** to commit your changes. In order words, you are not required to step through all screen to reach the **Activation & Summary** screen before you can save your changes.



Note: The **Save** button is available on most screen. If a screen does not show a **Save** button, click **Next** or **Done** until you reach to a screen where you can use its **Save** button to commit your changes.

Access SP connections

The **IdP Configuration** menu displays a list of the most-recently modified SP connections. You may create or import a connection. You may also edit a recently modified connection by clicking on its connection name.

To access the rest of the connections, click **Manage All** to open the **SP Connections** screen.

SP Connections

The **SP Connections** screen displays 20 connections at a time. As needed, you can use the pagination controls to navigate through the rest of your connections. You can also search connections by their names or connection IDs.



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.

You can sort by connection name, partner's connection ID, and default virtual server ID; narrow by protocol and status; and perform the following actions.

Create a connection

Follow the **Connection** wizard to create a new connection to your SP partner.

Copy a connection

Follow the **Connection** wizard to create a new connection based on an existing (source) connection. This is most useful if the new connection and the source connection have many setting values in common. Note that the new connection stays inactive until you change its status.

Export a connection

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

Follow the **Import Connection** wizard to import a connection export. If the connection already exists, you have the option to overwrite the existing connection.



Note: Prior to the import, you may modify the XML file to suit your needs. The XML file may also be imported to another PingFederate® environment acting in the same federation role (IdP or SP) at your site. Note that the source and the target must run the same version of PingFederate.

Export SAML metadata

Follow the **Export Metadata** wizard to generate a metadata for any SAML Browser SSO connection.

Update a SAML connection with metadata

Follow the **Import Metadata** wizard to update a SAML Browser SSO connection with metadata. You may update a connection via a metadata XML file or a metadata URL.



Important: The update operation may require additional configuration. Reviewing the connection after the update operation is recommended.

Delete a connection

Delete an inactive connection. You must click **Save** on the **SP Connections** screen to confirm your request. If you change your mind, click **Undelete** to cancel your request.

Override logging mode

Override the verbosity of runtime transaction logging for all SP connections.

Edit a connection

Open the connection by clicking on its name, select the setting that you want to reconfigure, and complete the change.

Resolve SP connection errors

PingFederate® automatically validates configured connections before displaying them on the **Connections** screen. 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.

- To resolve the error, select the connection and follow the on-screen instructions to modify the configuration, one connection at a time.

Import a connection

Use the **Import Connection** screen to import a connection.



Note: Prior to the import, you may modify the XML file to suit your needs. The XML file may also be imported to another PingFederate® environment acting in the same federation role (IdP or SP) at your site. Note that the source and the target must run the same version of PingFederate.

1. Click **Import**.
2. In the **Import Connection** screen, browse to a connection XML file.
3. Optional: Select the **Allow Update** check box. When selected, if the connection already exists, it will be overwritten.
4. Click **Import** and **Done**.

Update a SAML connection using metadata

You can update an existing SAML connection using a metadata file or a metadata URL from your partner.



Note: This manual update is independent from the optional and per-connection automatic update feature.

1. Go to the **Manage All > Connections** screen.
2. Click **Update with Metadata** under Action for the applicable SAML connection.
3. Refer to the following steps to import or update metadata. Instructions vary depending on the medium of the metadata.

Metadata medium Steps**A metadata file**

1. In the **Import Metadata** screen, select the **File** option.
2. 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** screen, and then click **Next**.

3. In the **Metadata Summary** screen, review the signature information to evaluate the authenticity of the metadata.
4. Click **Next**.

A new metadata URL

1. On the **Import Metadata** screen, select the **URL** option.
2. Enter a new metadata URL.



Note: If you specify a URL that has already been entered into the system through the **Server Configuration > Metadata URLs** workflow, the administrative console reminds you to select it from the **Existing URL Name** list.

3. Click **Load Metadata**.



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** screen, and then click **Next**.

4. In the **Metadata Summary** screen, review the signature information to evaluate the authenticity of the metadata.
5. Click **Next**.

An existing metadata URL

1. On the **Import Metadata** screen, select the **URL** option.
2. Select an existing metadata from the list, and then verify the URL in the **Selected Existing URL**.
3. Click **Load Metadata**.



Note: If there is a digital signature error, correct the issue using the **Server Configuration > Metadata URLs** workflow.

4. Click **Next**.



Note: If the endpoints in the metadata share the same base URL (protocol, hostname, and port), PingFederate 7.3 uses this information to populate the Base URL field. Consequently, individual endpoints on other screens do not include this information; only relative paths are shown.

Example

An 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 IdP can update the SP connection using the metadata immediately.


1. Access the **IdP Configuration > Manage All > Connections** screen.
2. Click **Update with Metadata** under Action for the applicable SAML connection.
3. Follow the workflow to complete the task.

Choose an SP connection template

On the **Connection Template** screen (shown only for a new connection), you can choose a quick-connection template if your installation includes an optional PingFederate® SaaS Connector.


 **Tip:** When you select a connection template, many connection settings are configured for you automatically.


- To configure a connection without a template, click **Next**.
- To use a template, select that option, then choose the template and enter additional information as required.

 **Note:** Once you click **Next**, you cannot return to this screen and make a different selection. If you intended to use a different template or no template, you must create a new connection.

Choose an SP connection type

If you are not using a connection template (which preconfigures browser-based SSO), indicate on the Connection Type screen whether the connection to this partner is for Browser SSO, WS-Trust STS, outbound provisioning, or any combination of them.


 **Tip:** 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. Note that each connection must use a unique (partner) connection ID.

- To configure a connection for secure browser-based SSO, select the **Browser SSO Profiles** check box.

If you have selected multiple protocols on the **Server Configuration > Server Settings > Roles & Protocols** screen and 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, namely **SAML 1.1** 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.

- To configure an STS connection, select the **WS-Trust STS** check box.

The **WS-Trust STS** option is only available after you enable the **WS-Trust** role on the **Server Configuration > Server Settings > Roles & Protocols** screen.

- To configure a connection for outbound provisioning, make that selection and then select the provisioning type from the list.

The **Outbound Provisioning** option is only available after you enable the **Outbound Provisioning** role on the **Server Configuration > Server Settings > Roles & Protocols** screen.

- 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.

Choose SP connection options

On the **Connection Options** screen (shown only for browser-based SSO connections), you can enable browser-based SSO, Attribute Query, or both for the current connection. For browser-based SSO, you may also enable IdP Discovery.

- To create a connection for browser-based SSO, select the **Browser SSO** check box.
- To enable IdP Discovery for this connection, make that selection (after selecting the **Browser SSO** check box).

Note that the **IdP Discovery** check box is only available if IdP Discovery is configured (see [Configure standard IdP Discovery](#) on page 175).


- To create a connection to facilitate the SAML 2.0 Attribute Query profile, select the **Attribute Query** check box (see [Attribute Query and XASP](#) on page 43).

Import SP metadata

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


- Enter a new metadata URL
- Select an existing metadata URL

 **Tip:** Using a metadata URL streamlines the configuration process. For example, by importing a metadata URL, you can quickly establish a Browser SSO connection to an InCommon-participating partner (see www.incommon.org/participants).


When you enter a new URL or select an existing metadata URL, PingFederate® also enables the automatic update option and checks the metadata daily. If PingFederate detects changes in the partner's digital signing certificates, encryption key, or contact information, it updates the connection automatically. If you prefer to update the connection manually, you can clear the **Enable Automatic Reloading** check box.

The frequency is configurable through the **Reload Delay** field on the **Server Configuration > Server Settings > Metadata Lifetime** screen.






Although optional, it is recommended that you turn on email notifications for SAML metadata update events on the **Server Configuration > Server Settings > Runtime Notification** screen.



 **Note:** If the metadata contains changes that require additional configuration, the email message also provides a list of the applicable items.

After the connection is created, you can add, remove, or change the metadata URL associated with the connection in the **Import Metadata** screen. In addition, you can also toggle the **Enable Automatic Reloading** option for the connection.

 **Tip:** These capabilities also lower the cost of maintaining connections with InCommon participants.

Refer to the following steps to import or update metadata. Instructions vary depending on the medium of the metadata.

Metadata medium	Steps
A metadata file	<ol style="list-style-type: none"> 1. On the Import Metadata screen, select the File option. 2. Choose the metadata file, and then click Next. <ul style="list-style-type: none">  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 screen, and then click Next. 3. On the Metadata Summary screen, review the signature information to evaluate the authenticity of the metadata. 4. Click Next.
A new metadata URL	<ol style="list-style-type: none"> 1. On the Import Metadata screen, select the URL option. 2. Enter a new metadata URL. <ul style="list-style-type: none">  Note: If you specify a URL that has already been entered into the system through the Server Configuration > Metadata URLs workflow, the administrative console reminds you to select it from the Existing URL Name list. 3. Optionally, clear the Enable Automatic Reloading check box to disable automatic update. 4. Click Load Metadata. <ul style="list-style-type: none">  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

Metadata medium	Steps
	<p>verification certificate on the Import Certificate screen, and then click Next.</p> <ol style="list-style-type: none"> On the Metadata Summary screen, review the signature information to evaluate the authenticity of the metadata. Click Next.
An existing metadata URL	<ol style="list-style-type: none"> On the Import Metadata screen, select the URL option. Select an existing metadata from the list, and then verify the URL in the Selected Existing URL. Optionally, clear the Enable Automatic Reloading check box to disable automatic update. Click Load Metadata. <ul style="list-style-type: none">  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, correct the issue using the Server Configuration > Metadata URLs workflow. Click Next.

Identify the SP

On the **General Info** screen, 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.

Provide the basic information to identify your partner.

Refer to the following table for more information.

Field	Description
Partner's Entity ID, Issuer, Partner's Realm, or Connection ID (Required)	<p>The published, protocol-dependent, unique identifier of your partner.</p> <p>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 may have been obtained out-of-band or via a SAML metadata file.</p> <p>For a WS-Federation connection, this is your partner's Realm.</p> <p>For an STS-only connection, this ID can be any unique identifier.</p>
Connection Name (Required)	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.
Virtual Server IDs	<p>If you want to identify your server to this connection partner using an ID other than the one you specified on the Server Configuration > Server Settings > Federation Info screen, enter a virtual server ID in this field and click Add.</p> <p>Enter additional virtual server IDs as needed.</p>
Base URL	The fully qualified hostname and port on which your partner's federation deployment runs (for example, <code>https://www.example.com:9031</code>). 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.

Field	Description
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 (<code>IdpAuthenticationAdapterV2</code>) in the PingFederate Java SDK. (Note that 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 (<code>IdpAuthenticationAdapterV2</code>) 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.

Configure Browser SSO

Browser-based SSO (also known as Browser SSO) relies on a user's web browser and HTTP requests to broker identity-federation messaging (in XML or JWT) between an IdP and an SP (in contrast to WS-Trust STS messaging, which is typically application-driven across the back channel and does not require browser mediation).

- To continue, click **Configure Browser SSO**.

IdP Browser SSO configuration steps

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 (see *Federation roles* on page 30). 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 SSO configuration steps

- [Choose SAML 2.0 profiles](#) on page 290
- [Set an SSO token lifetime](#) on page 290
- [Configure SSO token creation](#) on page 291
 - [Choose an identity mapping method](#) on page 291
 - [Set up an attribute contract](#) on page 292
 - [Manage authentication source mappings](#) on page 294
- [Configure protocol settings](#) on page 301
 - [Set Assertion Consumer Service URLs \(SAML\)](#) on page 302
 - [Specify SLO service URLs \(SAML 2.0\)](#) on page 304
 - [Choose allowable SAML bindings \(SAML 2.0\)](#) on page 304
 - [Set an artifact lifetime \(SAML\)](#) on page 304
 - [Specify artifact resolver locations \(SAML 2.0\)](#) on page 305
 - [Define signature policy \(SAML\)](#) on page 305
 - [Specify XML encryption policy \(for SAML 2.0\)](#) on page 364

WS-Federation SSO configuration steps

- [Set an SSO token lifetime](#) on page 290
- [Configure SSO token creation](#) on page 291
- [Define signature policy \(SAML\)](#) on page 305

- [Choose an identity mapping method](#) on page 291
- [Set up an attribute contract](#) on page 292
- [Manage authentication source mappings](#) on page 294
- [Configure protocol settings](#) on page 301
 - [Define a service URL \(WS-Federation\)](#) on page 303


SAML 1.x SSO configuration steps

- [Set an SSO token lifetime](#) on page 290
- [Configure SSO token creation](#) on page 291
 - [Choose an identity mapping method](#) on page 291
 - [Set up an attribute contract](#) on page 292
 - [Manage authentication source mappings](#) on page 294
- [Configure protocol settings](#) on page 301
 - [Set Assertion Consumer Service URLs \(SAML\)](#) on page 302
 - [Set a default target URL \(SAML 1.x\)](#) on page 303
 - [Set an artifact lifetime \(SAML\)](#) on page 304
 - [Define signature policy \(SAML\)](#) on page 305

After configuring SSO settings, you will normally need to configure authentication credentials, the range of which depends on your SSO selections (see [Configure credentials](#) on page 310). Also, other configuration tasks may remain to be configured for new or modified connections, depending on the selected options on the **Connection Options** screen.

Choose SAML 2.0 profiles

On the **SAML Profiles** screen, select one or more SAML 2.0 profiles.

 **Tip:** A SAML profile is the message-interchange scenario that you and your federation partner have agreed to use (in contrast to SAML binding, which is the transport protocol of SAML messages).

Note that the **SAML Profiles** screen is not shown for SAML 1.x connections because IdP SSO is assumed, SLO profiles are not supported, and the server supports the "destination-first" (SP-initiated) profile SSO automatically. This screen is also not presented for WS-Federation connections because profile selection is not required.

For SAML 2.0, PingFederate® supports all IdP- and SP-initiated SSO and SLO profiles. (For information on typical SSO and SLO profile configurations, including illustrations, see [SAML 2.0 profiles](#) on page 34.)

- Select the applicable profile (or profiles) based on your partner agreement.

You must always select at least one SSO profile.

Set an SSO token lifetime

Identity-federation standards require a window of time during which an SSO token is considered valid. Each SSO token has an issuance time-stamp element and elements indicating the allowable lifetime of the SSO token (in minutes) before and after the issuance time stamp.

- 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.

Configure SSO token creation

As an IdP, you must specify how PingFederate® obtains user-authentication information and use it to create SSO tokens appropriate for your SP partner, including additional user attributes as needed.

If you are a federation hub, bridging a service provider to one or more identity providers, you may associate one or more authentication policy contracts to the SP connection (see [Federation hub use cases](#) on page 86 for more information).

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.

- To continue, click **Configure ...**

Choose an identity mapping method

On the **Identity Mapping** screen, you choose the type of name identifier (*Name ID*) your partner requires. Your selection may affect the way that the 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 which protocol you are using.


- For information about SAML selections, see [Select a SAML Name ID type](#) on page 291.
- For information about WS-Federation selections, see [Select a WS-Federation Name ID type](#) on page 292.

Note that the **Identity Mapping** screen is not applicable to connections using the WS-Federation protocol in conjunction with 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.

Select a SAML Name ID type

The choice you make on the **Identity Mapping** screen affects how your SP partner makes use of account mapping or account linking.

If your SP is using account linking, establishing an attribute contract is not required. However, depending on your agreement, you may 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 (see [About attributes](#) on page 70).

 **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.

Select the type of name identifier that you and your SP have agreed to use.

Option	Description
Standard	<p>Select the Standard option if you want to send a known attribute to identify a user (for example, a username or an email address).</p> <p>In this scenario, the SP often uses account mapping to identify the user locally.</p>
Pseudonym	<p>Select the Pseudonym option 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.</p> <p>The identifier may also be used by the SP to make a persistent association between the user and a specific local account (account linking).</p> <p>Select the Include attributes ... check box if you want to set up an attribute contract to use in conjunction with an opaque identifier.</p>
Transient	<p>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.</p>

Option	Description
	A typical application for this selection might be, for example, when an SP provides generalized group accounts based on organizational rather than individual identity.
	Select the Include attributes ... check box if you want to set up an attribute contract to use in conjunction with an opaque identifier.

Select a WS-Federation Name ID type

On the **Identity Mapping** screen, indicate to the nature of the subject identifier by selecting one the three Name ID types for WS-Federation. Your selection may affect the way that the SP looks up and associates your users to their local accounts.

Note that the **Identity Mapping** screen is not applicable to connections using the WS-Federation protocol in conjunction with 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.

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 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 data store.
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 data store.
Common Name	This selection provides for anonymous SSO to your SP, generally using a hard-coded generalized logon. Make this selection if your partner agreement involves a many-to-one use case; for instance, if the SP has a group account set up for all users in a particular domain.


Set up an attribute contract

An attribute contract is the set of user attributes that you and your partner have agreed will be sent in the SSO tokens for this connection (see [Attribute contracts](#) on page 71). You identify these attributes on the **Attribute Contract** screen.

Establishing an attribute contract is required for WS-Federation connections. For SAML connections, attribute contracts are optional if you are sending either pseudonym or a transient identifiers to the partners. When establishing an attribute contract, you may 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** screen.


Protocol	Identity mapping	Attribute contract	SAML_SUBJECT	Additional attributes
SAML 2.0 or SAML 1.1	Standard	Required	Built-in. Subject name format can be changed by selecting a value from a list.	Optional. Attribute name format can be changed by selecting a value from a list.
SAML 2.0 or SAML 1.1	Pseudonym or Transient	Required <i>only</i> if the Include attributes ... check box is selected on the Identity Mapping screen; otherwise the	Assumed and cannot be added as an additional attribute.	At least one is required. Attribute name format can be changed by selecting a value from a list.

Protocol	Identity mapping	Attribute contract	SAML_SUBJECT	Additional attributes
		Attribute Contract screen is not shown.		
SAML 1.0	Standard	Required	Built-in. Subject name format can be changed by selecting a value from a list.	Optional. There is no attribute name format.
SAML 1.0	Pseudonym or Transient	Required <i>only if the Include attributes ... check box is selected on the Identity Mapping screen; otherwise the Attribute Contract screen is not shown.</i>	Assumed and cannot be added as an additional attribute.	At least one is required. There is no attribute name format.
WS-Federation in conjunction with SAML 1.1 as the token type	Email address, user principal name, or common name	Required	Built-in. There is no subject name format.	Optional. Attribute name format can be changed by selecting a value from a list.
WS-Federation in conjunction with JWT as the token type	Not applicable	Required	Not applicable	At least one is required. There is no attribute name format.

 **Tip:** If you are creating or updating a SAML 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.

1. Optional: Select a different name format for the built-in subject identifier, **SAML_SUBJECT**.

Applicable if you and the SP have agreed to a specific format (see [Attribute contracts](#) on page 71).

 **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 of PingFederate is required to activate any changes made to this file.)

(Other conditions described in the table apply.)

2. Extend the contract with additional attributes.

(Conditions described in the table apply.)

a) 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.


 **Tip:** 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 IdP application.

The value of this attribute can then be mapped later on the **Attribute Contract Fulfillment** screen (see [Configure attribute contract fulfillment](#) on page 297). Note that 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:

- `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

 **Tip:** If you are configuring a WS-Federation connection to Microsoft Windows Azure Pack, add `upn` to the JWT's attribute contract.


 **Tip:** If you are configuring a SAML connection to an InCommon participant (see www.incommon.org/participants), the attribute contract may 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](http://tools.ietf.org/html/rfc4524) (tools.ietf.org/html/rfc4524) and [RFC4519](http://tools.ietf.org/html/rfc4519) (tools.ietf.org/html/rfc4519). The following table describes a subset of the OIDs (object IDs) referenced by the most common attributes used by InCommon participants.

OID value	Description
0.9.2342.19200300.100.1.3	mail
1.3.6.1.4.1.5923.1.1.1.6	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	o
2.5.4.42	givenName
2.16.840.1.113730.3.1.241	displayName

For other attributes, refer to 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.

b) Select an attribute name format from the list.

Applicable if you and the SP have agreed to a specific format (see [Attribute contracts](#) on page 71).

 **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 of PingFederate is required to activate any changes made to this file.)

c) Click **Add**.

d) Repeat until all desired attributes are defined.


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.

Manage authentication source mappings


IdP adapters are responsible for handling user authentication as part of an SSO operation. A configured adapter in PingFederate® is known as an adapter instance.

In a basic scenario, you map an IdP adapter instance to an SP connection on the **Authentication Source Mapping** screen 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 may map multiple APCs to an SP connection (because your policies use multiple APCs), the same APC to multiple SP connections (because you want to reuse authentication policies on multiple connections), or a combination of them.

 **Tip:** To learn more about authentication policies and contracts, see [Authentication policies](#) on page 209.

Furthermore, you can 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:** To learn more about federation hub, see [Federation hub use cases](#) on page 86.

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 have the option to leave the decision to the users or create authentication policies to mandate authentication requirements.) Because each adapter instance or APC may return different user attributes, each mapping must define how the attribute contract is fulfilled in its mapping configuration.

- To map an IdP adapter instance, click **Map New Adapter Instance**.
- To map an APC, click **Map New Authentication Policy**.
- To edit the mapping configuration of an IdP adapter instance or APC, open it by clicking on its name, select the setting that you want to reconfigure, and complete the change.
- To remove an IdP 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 authentication sources, click **Done**.
- If you are editing an existing configuration and want to keep your changes, click **Save**.

When authentication sources (IdP adapter instances or connection mapping contracts) are restricted to certain virtual server IDs, the allowed IDs are displayed under **Virtual Server IDs**.

Select an authentication source


The first task of the mapping configuration is to map an adapter instance or an authentication policy contract to your connection.

- To map an adapter instance, follow these steps:
 - a) Select an adapter instance from the list.

If you do not see the desired adapter instance, click **Manage Adapter Instances** to create a new instance of any deployed adapter.

- b) Select the **Override Instance Settings** check box if you want to customize one or more adapter settings for this connection alone.

When selected, the administrative console adds a new set of sub tasks (the **Override Instance** screen and its sub tasks).

 **Tip:** 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 plug-in configurations](#) on page 69.

- To map an APC, select an adapter instance from the list.


If you do not see the desired APC, click **Manage Authentication Policy Contracts** to create a new policy contract.

If you are editing a currently mapped adapter 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 adapter settings specifically for this connection.

If you are editing a currently mapped APC, no changes can be made on this screen.


Override an IdP adapter instance

On the **Override Instance** screen, 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.

- Click **Override Instance Settings**.

On each of the settings screens, select the **Override** check box, make your changes, and then click **Next**. When you are finished, click **Done** to continue with the rest of the mapping configuration.

 **Note:** The override setting screens are functionally identical to those used for creating a new adapter instance (see [Manage IdP adapters](#) on page 275).

If you are editing a currently mapped adapter instance, click **Override Instance Settings** to reconfigure any overridden settings for this connection. You may also remove all overridden settings on a per-screen basis by clearing the **Override** check box near the top of the screen.

Restrict an authentication source to certain virtual server IDs

When you multiplex one connection for multiple environments (see [Multiple virtual server IDs](#) on page 84), you can enforce authentication requirements by restricting an authentication source to certain virtual server IDs on the **Virtual Server IDs** screen. By default, no restriction is imposed.


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.

If you are editing a currently mapped adapter instance or APC, you can toggle the **Restrict Virtual Server IDs** setting. You can also change the allowed virtual server IDs.

Select an attribute mapping method


On the **Mapping Method** screen, you select if and how PingFederate® should query local data stores to help fulfill the attribute contract in conjunction with attribute values from the authentication source.

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, determine whether local data stores can supply it.

 **Tip:** Alternatively, you may configure data store queries as part of the fulfillment configuration for the applicable IdP adapter contract or authentication policy contract. If so, you do not need to set up data store query on the connection level.

For more information, see [Define the IdP adapter contract](#) on page 277 or [Apply policy contracts to authentication policies](#) on page 221.

- Select the **Retrieve additional attributes from multiple data stores using one mapping** option if you want to configure one or more data stores to look up attribute for a single mapping.
- Select the **Retrieve additional attributes from a data store—including options to use alternate data stores and/or a failsafe mapping** option if you want to define alternate data stores to look up attribute and a failsafe mapping configuration.

 **Note:** When selected, the token authorization framework (through issuance criteria) does not apply. For more information, see [About token authorization](#) on page 74 and [Attribute mapping with multiple data sources](#) on page 506.

- Select the **Use only ...** option if you do *not* require connection-level data store query.

If you are editing a currently mapped adapter instance or APC, you can change the mapping method, which may require additional configuration changes in subsequent tasks.

Configure attribute sources and user lookup

Attribute sources are specific data store or directory locations containing information that may be needed for the attribute contract. You can use more than one attribute source when mapping values to the attribute contract. The order matters and affects the queries differently based on the selection made on the **Mapping Method** screen.

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—including options to use alternate data stores and/or a failsafe mapping

As soon as a query succeeds, PingFederate® moves on to the next task (contract fulfillment); therefore, you should prioritize the attribute sources.

- Click **Add Attribute Source** and then follow a series of sub tasks to complete the configuration.

For step-by-step instructions, see [Choose a data store](#) on page 508.

Repeat to add additional attribute sources.

If you are editing a currently mapped adapter instance or APC, you can add, remove, or reorder attribute sources, which may require additional configuration changes in subsequent tasks.

Configure attribute contract fulfillment

On the **Attribute Contract Fulfillment** screen, map values to the attributes defined for the contract. These are the values that will be included in the SSO tokens sent to the SP.

If you are bridging one or more identity providers to a service provider, map values to an authentication policy contract (see [Federation hub use cases](#) on page 86).

At runtime, an SSO operation fails if PingFederate® cannot fulfill the required attribute.

- For each attribute, select a source from the list and then choose or enter a value.


- **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 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 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 especially 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 87.

 **Note:** The **HTTP Request** context value is retrieved as a Java object rather than text. For this reason, OGNL expressions are more appropriate to evaluate and return values (see **Expression**).

- **LDAP, JDBC, or Custom**

When selected, the **Value** list is populated 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; for example, 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 (see **Text**).

Note that expressions are not enabled by default. For more information about enabling and editing OGNL expressions, see *Attribute mapping expressions* on page 497.

- **Text**

When selected, the text you enter is mapped to the value of the attribute in the SSO 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 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}
```

 **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).

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.

All attributes must be mapped.

If you are editing a currently mapped adapter instance or APC, you can update the mapping configuration, which may require additional configuration changes in subsequent tasks.

Map default attribute contract fulfillment

If you have chosen the failsafe option on the **Mapping Method** screen and the **Send user to SP using default list of attributes** option on the **Failsafe Attribute Source** screen, define the default values that should be sent in the SSO tokens to the SP on the **Attribute Contract Fulfillment** screen.


- For each attribute, select a source from the list and then choose or enter a value.
 - **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 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 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 especially 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 87.



Note: The **HTTP Request** context value is retrieved as a Java object rather than text. For this reason, OGNL expressions are more appropriate to evaluate and return values (see **Expression**).

- **Expression** (when enabled)

This option provides more complex mapping capabilities; for example, 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 (see **Text**).

Note that expressions are not enabled by default. For more information about enabling and editing OGNL expressions, see [Attribute mapping expressions](#) on page 497.

- **Text**

When selected, the text you enter is mapped to the value of the attribute in the SSO 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).

All attributes must be mapped.

If you are editing a currently mapped adapter instance or APC, you can update the mapping configuration, which may require additional configuration changes in subsequent tasks.

Specify issuance criteria

On the **Issuance Criteria** screen, you may define the criteria that PingFederate® can evaluate to determine whether to issue an SSO token for a user. This optional token authorization configuration can be used to restrict who can SSO to protected resources. (For more information, see [About token authorization](#) on page 74.)

Note that token authorization does not apply if you have chosen the failsafe option on the **Mapping Method** screen (and the **Issuance Criteria** screen is not shown).

Issuance criteria and OGNL expressions

When defining multiple criteria using the user-interface elements, all criteria must pass in order for a user to be authorized. Source attribute values are evaluated against specified values using an exact match algorithm. The **multi-value contains** (or **multi-value does not contain**) conditions (intended for attributes that may contain multiple values) return true if the source attribute contains the exact specified value. For more complex evaluations, including conditional criteria and partial matching, define criteria using OGNL expressions.




Important: When you multiplex one connection for multiple environments (see [Multiple virtual server IDs](#) on page 84), consider using an OGNL expression to verify the virtual server ID in conjunction with other conditions, such as group membership information, to protect against unauthorized access (see [Sample OGNL expressions](#) on page 498).

Note that expressions are not enabled by default. For more information about enabling and editing OGNL expressions, see [Attribute mapping expressions](#) on page 497.

1. Select the source of the attribute that you want PingFederate to evaluate from the **Source** list.


When you select a source, the **Attribute Name** list is populated with the associated attributes or properties, as described in the following table.

Source	Attribute name
Adapter or Authentication Policy Contract (the authentication source)	Populated with the attributes from the authentication source.
Context	Populated with the available context of the transaction.  Note: The HTTP Request context value is retrieved as a Java object rather than text. For this reason, OGNL expressions are more appropriate to evaluate and return values.
LDAP, JDBC, or Custom (if a data store is configured)	Populated with attributes that you have selected in the attribute source configuration.
Mapped Attributes	Populated with attributes from the attribute contract.

2. Select an attribute from the **Attribute Name** list.
3. Select the evaluation method from the **Condition** 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

 **Note:** Use one of the *multi-value* conditions when you want PingFederate to verify that an attribute contains or does not contain the specified value in its attribute list. Using a single-value condition when an attribute has multiple values causes the condition to evaluate consistently to false.

4. Enter the exact value in the **Value** field.

Wildcards (for partial matching) are not allowed.

 **Tip:** You can use OGNL expressions to perform more complex evaluations, including partial matching.

5. Optional: Use the **Error Result** field to customize an error code or error message for use if the token authorization fails.


Error results are handled differently for SP-initiated SSO and IdP-initiated SSO.

SP-Initiated SSO

- **SAML:** The **Error Result** field value is used by the StatusMessage element in the response to the SP.
- **Template (WS-Federation):** The **Error Result** field value is used by the *\$errorDetail* variable in the `sourceid-wsfed-idp-exception-template.html` template file.

IdP-initiated SSO


- **Redirect:** When an InErrorResource URL is provided, the **Error Result** field value is used by the ErrorDetail query parameter in the redirect URL.
- **Template:** When an InErrorResource URL is not provided, the **Error Result** field value is used by the *\$errorDetail* variable in the `idp.sso.error.page.template.html` template file.

-  **Tip:** You may also enter an error code in the **Error Result** field that the error template or the web application can use to process the error result in a variety of ways; for example, a redirect to another page or e-mail to an administrator.

If you leave this field blank, a default `ACCESS_DENIED` error result is used if the authorization fails.


6. Click **Add**.

7. Repeat these steps if you want to add additional criteria.


-  **Important:** All criteria must pass in order for a user to be authorized.

8. Optional: Click **Show Advanced Criteria** to configure more complex evaluations using OGNL expressions.

- a) Enter an OGNL expression in the **Expression** field.
- b) Enter an error code or an error message in the **Error Result** field.

 **Note:** If the OGNL expression resolves to a string value (rather than `true` or `false`), the returned value overrides the **Error Result** field value.

- c) Click **Add**.
- d) Click **Test**, enter values in the applicable fields, and test the resulting output for the expression.
- e) Repeat these steps if you want to add additional criteria using expressions.

-  **Important:** All criteria must pass in order for a user to be authorized.

If you want conditional criteria, add only one criterion with such processing logic.

If you are editing a currently mapped adapter instance or APC, you can add, remove, or edit any issuance criteria.

Review the authentication source mapping

When you have finished adding a new or modifying an existing instance of an adapter or an authentication policy contract, you can review the configuration on its **Summary** screen.

- If you need to make any changes, click the heading over the information you want to edit.
- If you are creating a new connection and want to keep your configuration, click **Done**.
- If you are editing an existing configuration and want to keep your changes, click **Save**.

Review the SSO token creation summary

When you are finished with the **Assertion Creation** task, you can review the configuration on its **Summary** screen.

- If you need to make any changes, click the heading over the information you want to edit.
- If you are creating a new connection and want to keep your configuration, click **Done**.
- If you are editing an existing configuration and want to keep your changes, click **Save**.

Configure protocol settings

The **Protocol Settings** screen provides the launching point for configuring partner endpoints, message customizations, and other protocol-specific settings for browser-based SSO connections.

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
- To continue, click **Configure Protocol Settings**.

Set Assertion Consumer Service URLs (SAML)

The assertion consumer service (ACS) endpoint is a location to which the SSO tokens are sent, according to partner requirements. ACS is applicable to all SAML versions and both the IdP- and SP-initiated SSO profiles.

On the **Assertion Consumer Service URL** screen, select the applicable SAML binding and enter the corresponding ACS endpoint URL.



Note: The SP may request that the SAML assertion be sent to one of several URLs, via 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 may 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.

1. Configure one or more SAML ACS endpoints.

- Select a SAML binding from the list; for example, **POST**.
- Enter the ACS endpoint URL to the **Endpoint URL** field.

You may enter a relative path (begin with a forward slash) if you have provided a base URL on the **General Info** screen.

- Make the selection 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 may reset the default selection when you add another ACS endpoint.

- 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.

- Click **Add**.
- Optional: Repeat to add additional ACS endpoints.

2. Optional: Customize messages using OGNL expressions.

Note that expressions are not enabled by default. For more information about enabling and editing OGNL expressions, see [Attribute mapping expressions](#) on page 497.

- Click **Show Advanced Customizations**.
- Select a message type from the list.
- Enter an OGNL expression to fulfill your use case.



Note: For more information about **Message Type**, available variables, and sample OGNL expressions, see [Customize assertions and authentication requests](#) on page 501 .

- Click **Add**.
- Optional: Repeat to add another message customization.

If you are editing an existing connection, you can reconfigure any items, which may require additional configuration changes in subsequent tasks. You must always configure at least one ACS endpoint.

Set a default target URL (SAML 1.x)

SAML 1.x SP connections requires that a default target URL be specified for a scenario where the IdP application does not include one in its SSO request. This default URL represents the destination on the SP where the user will be directed.

- Enter default destination in the **Default Target URL** field.

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.

Define a service URL (WS-Federation)

On the **Service URL** screen, enter the WS-Federation protocol endpoint of your SP partner where PingFederate® sends SSO tokens and SLO cleanup messages. The SSO tokens are transmitted within an RSTR (Request for Security Token Response) message in response to a request for authentication from the SP. SLO cleanup messages are sent to your partner when PingFederate (the IdP) receives a user's SLO request. Such 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 screen. If the option to validate wreply for SLO is enabled, these additional domains and paths will also be taken into consideration as well (see [Manage partner redirect validation](#) on page 177).


Some federation use cases may 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.

1. Enter the WS-Federation protocol endpoint at the SP site in the **Endpoint URL** field.

You may enter a relative path (begin with a forward slash) if you have provided a base URL on the **General Info** screen.

2. Specify additional allowed domains and paths.


- a) Optional: Indicate whether to 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.

- b) Enter the expected domain or IP address under **Valid Domain Name**.


Use the domain only, without qualifiers; for example: `example.com`

* `.example.com` covers `hr.example.com` or `email.example.com` but not `example.com`.

 **Important:** While using an initial wildcard provides the convenience of allowing multiple sub domains with one entry, consider adding individual sub domains to limit the redirection to a list of known hosts.

- c) Optional: Enter the exact (and case-sensitive) path of the resource under **Valid Path**.

Starts 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. For example, `/app/Consumer.jsp` allows `/app/Consumer.jsp` but rejects `/app/consumer.jsp`.

 **Tip:** You can also enter multiple query parameters with or without a fragment.

For instance, `/app/Consumer.jsp?area=West&team=IT#ref1001` matches `/app/Consumer.jsp?area=West&team=IT#ref1001` but not `/app/Consumer.jsp?area=East&team=IT#ref1001`.

- d) Optional: Select the check box under **Allow Any Query/Fragment** if you want to allow any query parameters or fragment in the resource.

 **Note:** When selected, no query parameter, fragment, or a combination of them is allowed in the path.

- e) Click **Add**.

- f) Optional: Repeat the previous step to add additional entries as needed.

Note that 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.

Note that expressions are not enabled by default. For more information about enabling and editing OGNL expressions, see *Attribute mapping expressions* on page 497.

- a) Click **Show Advanced Customizations**.
- b) 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 *Customize assertions and authentication requests* on page 501 .

- d) Click **Add**.
- e) Optional: Repeat to add another message customization.

If you are editing an existing connection, you can reconfigure any items, which may require additional configuration changes in subsequent tasks.

Specify SLO service URLs (SAML 2.0)

On the **SLO Service URLs** screen, you associate bindings to the endpoints where your SP receives logout requests when SLO is initiated at your site and where PingFederate® sends SLO responses when it receives SLO requests from the SP.

This step applies only to SAML 2.0 connections when either SLO profile is selected on the **SAML Profiles** screen.

1. Select a SAML binding from the list; for example, **POST**.
2. Enter the SLO endpoint URL to the **Endpoint URL** field.

You may enter a relative path (begin with a forward slash) if you have provided a base URL on the **General Info** screen.

3. Optional: Enter an URL to the **Response URL** field.

When specified, it 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 may enter a relative path (begin with a forward slash) if you have provided a base URL on the **General Info** screen.

4. Click **Add**.
5. Optional: Repeat to add additional SLO endpoints.

If you are editing an existing connection, you can reconfigure the SLO endpoints, which may require additional configuration changes in subsequent tasks.

Choose allowable SAML bindings (SAML 2.0)

On the **Allowable SAML Bindings** screen, you select the one or more bindings that your SP partner can use to send SAML authentication requests or SLO messages.

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** screen.

- Select only the applicable SAML binding (or bindings) based on your partner agreement.

If you have specified an ACS or SLO endpoint using the artifact (outbound) binding, you must including SOAP as one of the allowable (inbound) binding.

If you are editing an existing connection, you can reconfigure the allowable bindings, which may require additional configuration changes in subsequent tasks.

Set an artifact lifetime (SAML)

When PingFederate® sends an artifact to your 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** screen, specify the expiry information in seconds.

You can change the default value per your requirements, if needed. Also consider synchronizing clocks between your server and your partner's SAML gateway server. If clocks are not synchronized, you might need to set the artifact lifetime to a higher value.

- Optional: Override the default value of the **Artifact Lifetime** field.

The default value is 60 (seconds).

If you are editing an existing connection, you can update the artifact lifetime.

Specify artifact resolver locations (SAML 2.0)

When the artifact binding is enabled as one of the allowable bindings on the **Allowable SAML Bindings** screen, you must provide at least one artifact resolution service (ARS) endpoint on the **Artifact Resolver Locations** screen. This is the location where PingFederate® sends back-channel requests to resolve artifacts received from the SP.

1. Enter the ARS endpoint URL to the **URL** field.


You may enter a relative path (begin with a forward slash) if you have provided a base URL on the **General Info** screen.

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, starting from 0. If you want to define your own index values, you must make sure the index values are unique.

3. 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. Similarly, if one endpoint uses HTTP, then all must use HTTP.

If you are editing an existing connection, you may reconfigure any ARS endpoints.

Define signature policy (SAML)

The **Signature Policy** screen provides options controlling how digital signatures are used for SSO messages. The choices made on this screen depend on your partner agreement (see [Digital signing policy coordination](#) on page 77).

Digital signing is required for SAML response messages sent from the IdP via 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 **Always sign the SAML Assertion** option on this screen. Note that when the latter is selected, only the assertion portion of the SAMLResponse message is signed, not the entire SAML response message. (This is the only option that appears for SAML 1.x connections.)

SAML 2.0 authentication requests from the SP may 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** screen.)

- To continue, select the option (or options) based on your partner agreement.

If you are editing an existing connection, you can reconfigure the digital signature policy, which may require additional configuration changes in subsequent tasks.

Configure XML encryption policy (SAML 2.0)

For SAML 2.0 configurations, in addition to using signed assertions to ensure authenticity, you and your partner may also agree to encrypt all or part of an assertion to improve privacy. If so, you can configure these settings on the **Encryption Policy** screen. There are three options; each option provides different capabilities.

Option	Name identifier (SAML_SUBJECT)	Other attributes	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.

Option	Name identifier (SAML_SUBJECT)	Other attributes	Encrypt the SAML_SUBJECT in SLO messages to the SP	Allow encryption in SLO messages from the SP
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 <i>only if</i> you select to encrypt the name identifier (SAML_SUBJECT).	Available as an option <i>only if</i> you select to encrypt the name identifier (SAML_SUBJECT).

- To continue, select the option (and options) based on your partner agreement.

If you are editing an existing connection, you can reconfigure the XML encryption policy, which may require additional configuration changes in subsequent tasks.

Review protocol settings

When you are finished with the **Protocol Settings** task, you can review the configuration on its **Summary** screen.

- If you need to make any changes, click the heading over the information you want to edit.
- If you are creating a new connection and want to keep your configuration, click **Done**.
- If you are editing an existing configuration and want to keep your changes, click **Save**.

Review browser-based SSO settings

When you are finished with the **Browser SSO** primary task, you can review the configuration on its **Summary** screen.

- If you need to make any changes, click the heading over the information you want to edit.
- If you are creating a new connection and want to keep your configuration, click **Done**.
- If you are editing an existing configuration and want to keep your changes, click **Save**.

Configure the Attribute Query profile

At the Attribute Query step you configure your connection to respond to requests for user attributes from your partner SP, if you have chosen this option (see [Choose SP connection options](#) on page 286). Attribute queries are not dependent on single sign-on but may be used independently or in conjunction with Browser SSO or provisioning to provide flexibility in how a user authenticates with SP applications (see [Attribute Query and XASP](#) in the “Supported Standards” chapter of the *Get started with PingFederate* guide).

- To continue, click **Configure Attribute Query Profile**.

Define retrievable attributes

On this screen you specify the user attributes you and your partner have agreed to allow in an attribute query transaction. Note that the SP may 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.

To add an attribute:

- Enter the attribute name in the text box and click **Add**.

To edit an attribute name:

1. Click **Edit** and make your change.
2. Click **Update**.

To delete an attribute:

- Click **Delete**.

Configure attribute lookup

Attribute sources are specific data store or directory locations containing information that may be returned to the SP in response to an attribute request.

This portion of the attribute query configuration allows you to configure one or more data stores to look up attributes and to set up search parameters.

To configure an attribute source:

- Click **Add Attribute Source** and complete the setup steps (see [Choose a data store for attribute query](#) on page 307).

To modify an attribute source configuration:

1. Click the attribute source Description link.
2. Click **Save** on the screen you change.



Note: Depending on what you change, you may need to modify dependent data in subsequent steps, as indicated. Click **Save** or **Done** when either of those options appears.

To reach this screen for editing:

1. Click **IdP Configuration** on the Main Menu.
2. Click the connection name under SP Connections.
Click **Manage All**, if needed, to see a full list of connections.
3. Click **Attribute Query** under the SP Connection tab.
4. Click **Configure Attribute Query Profile**.
5. Click **Attribute Source & User Lookup** on the Summary screen.

Choose a data store for attribute query

The process of configuring PingFederate to look up attributes in a data store for attribute-query responses is similar to that used for SSO Attribute Sources and User Lookup. On the **Data Store** screen, choose a data store instance for PingFederate to look up attributes.

1. Enter a description (and ID if prompted) for the data store.
2. Select a data store instance from the **Active Data Store** list.



Tip: 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.

3. Depending on the data store type, the rest of the setup varies as follows:

Data store type	Required tasks
JDBC	<ul style="list-style-type: none"> • Specify a JDBC database table and columns on page 508 • Enter a database search filter on page 509
LDAP	<ul style="list-style-type: none"> • Specify an LDAP directory and attributes on page 510 • Define encoding for LDAP binary attributes on page 511 (optional) • Enter an LDAP search filter on page 511
Custom	<ul style="list-style-type: none"> • Specify a custom source filter and fields on page 512



Important: When attribute queries are sent using XASP, use the variable `SubjectDN`—rather than `SAML_SUBJECT`—to retrieve the subject identifier. You may also use any of these DN-parsing variables: `CN`, `OU`, `O`, `L`, `S`, `C`, and `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](#) in the “Supported Standards” chapter of the *Get started with PingFederate* guide.

Configure attribute mapping fulfillment


The last step in configuring an attribute source is to map values into the assertion to be sent in response to an attribute query.

You map attributes on the Attribute Mapping Fulfillment screen.


Map each attribute into the assertion from one of these Sources:

- **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 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 especially 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 87.

 **Note:** The **HTTP Request** context value is retrieved as a Java object rather than text. For this reason, OGNL expressions are more appropriate to evaluate and return values (see **Expression**).

- **LDAP/JDBC/Custom**

Values are returned from your attribute source. When you make this selection, the Value list is populated by the LDAP, JDBC, or Custom attributes you identified for this Attribute Source.

 **Note:** PingFederate appends a description in parentheses for each data store lookup (see [Choose a data store for attribute query](#) on page 307).

- **Expression (when enabled)**

This option provides more complex mapping capabilities—for example, transforming incoming values into different formats (see [Attribute mapping expressions](#) on page 497). All of the variables available for text entries (see below) are also available for expressions.

- **Text**

This can be text only, or you can mix text with references to any of the values from your user-data store using this syntax:

```
${ds.attr-source-id.attribute}
```

where `attr-source-id` is the **Attribute Source Id** value (see [Choose a data store for attribute query](#) on page 307) and `attribute` is any of the data store 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.

Specify issuance criteria for attribute query (optional)


Use this screen to define criteria PingFederate can evaluate to determine whether to issue an attribute query response (see [About token authorization](#) on page 74). This extension of token authorization can be used to restrict who can access protected resources.

To configure Issuance Criteria:


1. Select a source containing attributes for token authorization.

Associated attributes appear in the Attribute Name drop-down list:

- Context – Select to use values returned from the context of the transaction at runtime.

 **Note:** For the HTTP Request selection, because it is retrieved as an object rather than text, you must generally use OGNL expressions for evaluation. Click **Show Advanced Criteria** to use expression mapping. (If that button is not displayed on the screen, expressions are not enabled—see [Enable and disable expressions](#) on page 497). For syntax and examples, see sections under [Construct OGNL expressions](#) on page 498.)


- LDAP/JDBC/Custom (if a data store is configured) – Select to access attributes returned from a data source

 **Note:** PingFederate appends a description in parentheses for each data store lookup (see [Choose a data store for attribute query](#) on page 307).

- Mapped Attributes – Select to access retrievable attributes defined for the attribute query profile.


2. Select an attribute name.

3. Select the Condition you want to apply.

 **Note:** Use one of the *multi-value* conditions when you want PingFederate to verify that an attribute contains or does not contain the specified value in its attribute list. Using a single-value condition when an attribute has multiple values causes the condition to evaluate consistently to false.


4. Enter the exact value in the **Value** field.

Wildcards (for partial matching) are not allowed.

 **Tip:** You can use OGNL expressions to perform more complex evaluations, including partial matching.

5. Optional: Use the Error Result field to customize an error code or error message for use if the token authorization fails.


The Error Result field is used by the `StatusMessage` element in the SAML response to the SP.

 **Note:** 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.

If you leave this field blank, the user sees a default ACCESS_DENIED error result at runtime if the authorization fails.

6. Click **Add**.

7. Repeat the steps above to add additional attributes, as needed, for each authorization criteria.

 **Note:** All criteria must pass in order for a user to be authorized.


8. Optional: Click **Show Advanced Criteria** to configure more complex evaluations using OGNL expressions.

 **Note:** Expressions must be enabled for the Show Advanced Criteria button to appear (see [Enable and disable expressions](#) on page 497).


- Use the in-line editor box to enter the OGNL expression.

For more information about OGNL, see [Attribute mapping expressions](#) on page 497.

- Use the Error Result box to enter an error code or an error message for use if authorization fails (see step 5 above for more information)

 **Note:** If the OGNL expression resolves to a string value (rather than true or false), then the returned value overrides the Error Result.

- Click **Add**. Continue to add expressions, as needed.
- Click the **Test** link to input values and test the resulting output for the expression.

 **Note:** For more information on testing OGNL expressions, see [Using the OGNL edit screen](#) on page 501

9. Click **Next** (or **Done**, if you are finished with the configuration, and then **Save** on the Manage Mappings screen).

To edit Issuance Criteria:

- Click **Edit** under Actions for the criteria.

To delete Issuance Criteria:

- Click **Delete** under Actions for the criteria and then click **Save**.

Specify security policy

This screen allows you to specify the digital signing and encryption policy to which you and your partner have agreed. These selections will trigger requirements for setting up Credentials (see [Configure credentials](#) on page 310).

To configure attribute-query security policy for this partner:

- Check or clear the check boxes and click **Next** or **Done**.

Review the Attribute Query configuration

To reconfigure saved profiles:

1. Click the heading over the information you want to change.
2. Click **Done** on the screen containing your change.

If you need to make additional changes, do so and continue by clicking **Done** until you reach the Attribute Query screen.

3. Click **Save** on the Attribute Query screen.

Configure credentials

The **Credentials** screen 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. Your connection configuration may involve any or all of the following:

- [Configure back-channel authentication \(SAML\)](#) on page 310
- [Configure digital signature settings](#) on page 313
- [Configure signature verification settings \(SAML 2.0\)](#) on page 313
- [Select an encryption certificate](#) on page 314
- [Select a decryption key \(SAML 2.0\)](#) on page 315
- To continue, click **Configure Credentials**.

Configure back-channel authentication (SAML)

Depending on your Browser SSO use cases, the administrative console prompts you to configure authentication requirements for inbound messages, outbound messages, or both. Refer to the following table for more information.

Use case	Back-channel authentication requirements	Back-channel messages
A connection is configured with a SAML ACS endpoint that uses the artifact binding on the Protocol Settings > Assertion Consumer Service URL screen.	Inbound	Artifact resolution requests

Use case	Back-channel authentication requirements	Back-channel messages
A connection is configured with a SAML 2.0 SLO endpoint that uses the artifact binding on the Protocol Settings > SLO Service URLs screen	Inbound	Artifact resolution requests SOAP messages
A connection is configured with a SAML 2.0 SLO endpoint that uses the SOAP binding on the Protocol Settings > SLO Service URLs screen	Outbound	SOAP SLO messages
The SAML 2.0 Artifact binding is enabled on the Protocol Settings > Allowable SAML Bindings screen	Outbound	Outbound artifact resolution requests
The SOAP binding is enabled on the Protocol Settings > Allowable SAML Bindings screen	Inbound	Inbound SOAP messages
The SAML 2.0 Attribute Query profile is enabled on the Connection Options screen	Inbound	Inbound Attribute Query requests

Refer to the following topics for configuration steps:

- [Configure authentication requirements for outbound messages](#) on page 311
- [Configure authentication requirements for inbound messages](#) on page 312

Configure authentication requirements for outbound messages

1. On the **Back-Channel Authentication** screen, click **Configure** to the right of the list of messages under **Send to your partner**.
2. On the **Outbound SOAP Authentication Type** screen, 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)** screen.

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** screen. If you have not yet created or imported the client certificate, click **Manage Certificates** to do so (see [Manage SSL client keys and certificates](#) on page 182).

 **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 a subsequent screen, **Digital Signature Settings**.

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 may be used in any combination or independently.

3. On the **Summary** screen for outbound authentication requirements:


- If you need to make any changes, click the heading over the information you want to edit.
- If you are creating a new connection and want to keep your configuration, click **Done**.
- If you are editing an existing configuration and want to keep your changes, click **Save**.

Configure authentication requirements for inbound messages

1. On the **Back-Channel Authentication** screen, click **Configure** to the right of the list of messages under **Received from your partner**.
2. On the **Inbound Authentication Type** screen, 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)** screen.

-  **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 screens (see next step).

Digital Signature (Browser SSO profile only)

You select a signing certificate on a subsequent screen, **Signature Verification Settings**.

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 may 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 may 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** screen.

Anchored

The partner certificate must be signed by a trusted certificate authority (CA). Optionally, you may 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 **Server Configuration > Trusted CAs** screen.


Unanchored

The partner certificate is self-signed or you want to trust a specified certificate.

-  **Note:** When anchored certificates are used between partners, certificates may 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 77.)

Trust model	Subsequent steps
Anchored	<p>On the Subject DN screen:</p> <ol style="list-style-type: none"> 1. Enter the Subject DN of the certificate. 2. (Optional) Select the Restrict Issuer check box and enter the Issuer DN of the certificate.

Trust model	Subsequent steps
	 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 screen, 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 (see Manage certificates from partners on page 190).

4. On the **Summary** screen for inbound authentication requirements:

- If you need to make any changes, click the heading over the information you want to edit.
- If you are creating a new connection and want to keep your configuration, click **Done**.
- If you are editing an existing configuration and want to keep your changes, click **Save**.

Configure digital signature settings

Digital signing is required for browser-based SSO tokens and SLO messages sent via POST or redirect bindings. It is also required for WS-Trust STS SP connections (for the purpose of signing the outbound SAML security tokens).


On the **Digital Signature Settings** screen, select the certificate that you will use to sign the SSO tokens and SLO messages for this SP.

Note that, 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 the **Protocol Settings > Signature Policy** screen or the artifact resolution messages on the **Back-Channel Authentication > Outbound SOAP Authentication Type** screen.


If digital signing is not required, the **Digital Signature Settings** screen is not shown.

1. Select a signing certificate from the list.

If you have not yet created or imported your certificate into PingFederate, click **Manage Certificates** (see [Manage digital signing certificates and decryption keys](#) on page 184).

 **Note:** For WS-Federation connections using JSON Web Tokens, only EC and RSA certificates are supported. Furthermore, 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.

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.

 **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 JSON Web Tokens.

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.

Configure signature verification settings (SAML 2.0)

Depending on your partner agreement, digital signature processing may be required.


If you chose to require digital signatures on SAML 2.0 authentication requests on the **Protocol Settings > Signature Policy** screen or inbound messages on the **Back-Channel Authentication > Inbound Authentication Type** screen, you must configure the required certificate information that PingFederate[®] can use to verify the signed messages.

The **Signature Verification Settings** is the launching point for this task. If digital signature verification is not required, the **Signature Verification Settings** screen is not shown.

1. On the **Signature Verification Settings** screen, click **Manage Signature Verification Settings**.
2. On the **Trust Model** screen, select a trust model on the **Certificate Verification Method** screen.


Anchored

The partner certificate must be signed by a trusted certificate authority (CA). Optionally, you may 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 **Server Configuration > Trusted CAs** screen.



 **Important:** If you are using the redirect binding for SLO, you cannot use anchored certificates because SAML 2.0 does not permit certificates to be included using this transport method.

Unanchored

The partner certificate is self-signed or you want to trust a specified certificate.

 **Note:** When anchored certificates are used between partners, certificates may 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 77.)

Trust model	Subsequent steps
Anchored	<p>On the Subject DN screen:</p> <ol style="list-style-type: none"> 1. Enter the Subject DN of the certificate or extract it from your SP partner's certificate if the certificate is stored on an accessible file system. 2. (Optional) Select the Restrict Issuer check box and enter the Issuer DN of the certificate. Alternatively, extract it from your partner's certificate. <p> 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.</p>
Unanchored	<p>On the Signature Verification Certificate screen:</p> <ol style="list-style-type: none"> 1. Select a primary certificate from the list. <p>If you have not yet imported the certificate from your partner, click Manage Certificates to do so (see Manage certificates from partners on page 190).</p> <ol style="list-style-type: none"> 2. Select a secondary certificate from the list. <p> Tip: Use this field if your partner has sent you a new certificate to replace one that is ready to expire. The server will automatically verify against the secondary certificate when the primary one expires.</p>

3. On the **Summary** screen for signature verification:

- If you need to make any changes, click the heading over the information you want to edit.
- If you are creating a new connection and want to keep your configuration, click **Done**.
- If you are editing an existing configuration and want to keep your changes, click **Save**.

Select an encryption certificate

For browser-based SSO, if you chose to encrypt all or part of an SSO assertion on the **Protocol Settings > Encryption Policy** screen, you must identify the certificate that PingFederate® can use to do so.

You must also select a certificate if your requirements include encrypting an assertion in response to an attribute query on the **Attribute Query > Security Policy** screen.


For WS-Trust STS, this configuration is also required if you have enabled the **Generate Key for SAML Holder of Key Subject Confirmation Method** or **Encrypt SAML 2.0 Assertion** option (or both options) on the **WS-Trust > Protocol Settings** screen.

If encryption is not required, the **Select XML Encryption Certificate** screen is not shown.

1. Optional: Change the default settings under **Block Encryption Algorithm**.

Due to import control restrictions, the standard JRE distribution supports strong but not unlimited encryption. To use the strongest AES encryption, when permissible, download and install the appropriate version of "Java Cryptography Extension (JCE) Unlimited Strength Jurisdiction Policy Files" from the [Oracle Downloads](http://www.oracle.com/technetwork/java/javase/downloads/index.html) website (www.oracle.com/technetwork/java/javase/downloads/index.html).

For more information about XML block encryption and key transport algorithms, see [XML Encryption Syntax and Processing from W3C](http://www.w3.org/TR/xmlenc-core/) (www.w3.org/TR/xmlenc-core/)

 **Note:** Under **Key Transport Algorithm**, **RSA-v1.5** is disabled for new connections for security reasons. If you are updating an existing connection that uses RSA-v1.5, consider changing the selection for increased security.

Note that the JCE provider libraries distributed with SafeNet Luna SA HSM does not support the RSA-OAEP algorithm at the time of the PingFederate 7.1 R2 release. As a result, RSA-v1.5 is still allowed when deploying PingFederate with a Luna SA HSM.

2. Select a partner certificate from the list.

If you have not yet imported the certificate from your partner, click **Manage Certificates** to do so (see [Manage certificates from partners](#) on page 190).

Select a decryption key (SAML 2.0)

When you chose to encrypt the name identifier (SAML_SUBJECT) on the **Protocol Settings > Encryption Policy** screen, you also have the option to allow the SP to encrypt the name identifier in its SLO requests (if the SP-initiated SSO profile is enabled for the connection). To enable this inbound encryption, you must specify at least one certificate on the **Select Decryption Keys** screen.

If decryption is not required, the **Select Decryption Keys** screen is not shown.

1. Select the primary XML decryption key from the list.

If you have not yet created or imported your certificate into PingFederate, click **Manage Certificates** (see [Manage digital signing certificates and decryption keys](#) on page 184).

2. Optional: Select the secondary XML decryption key from the list.


Review credential settings

When you finish configuring various settings under the **Credentials** task, you can review the configuration on its **Summary** screen.

- If you need to make any changes, click the heading over the information you want to edit.
- If you are creating a new connection and want to keep your configuration, click **Done**.
- If you are editing an existing configuration and want to keep your changes, click **Save**.

Configure outbound provisioning

PingFederate[®]'s Outbound Provisioning allows an IdP to create and maintain user accounts at standards-based partner sites using SCIM as well as select-proprietary provisioning partner sites that are protocol-enabled (see [Outbound provisioning for IdPs](#) on page 81).

 **Note:** This configuration task is presented in the administrative console only when you enable Outbound Provisioning (see [Choose an SP connection type](#) on page 286).

- To continue, click **Configure Provisioning**.



Note: Screen illustrations in this section are presented in most cases for service providers using the SCIM protocol. However, the screens are comparable for SaaS vendors for whom provisioning is supported, and the information and instructions provided are the same unless otherwise noted.

Define a provisioning target

Information on the **Target** screen indicates 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.



Note: The target configuration settings vary among SCIM outbound provisioning and various SaaS provisioning.

Before proceeding further, obtain the users resource URL from your service provider; for example, <https://example.com/v1/Users>. If supported by the provider, also obtain the groups resource URL; for example, <https://example.com/v1/Groups>. Also determine whether the provider supports PATCH updates (see [step 5](#)) and prefers full distinguished name (DN) for group provisioning (see [step 6](#)).

For SCIM provisioning to PingOne®, sign on to the [PingOne admin portal](#) and review the target information on the **Setup > Identity Repository** tab.

For any SaaS Connector target, please refer to documentation in the add-on distribution package.

The following steps describe the fields required for the bundled PingFederate provisioning plug-in for SCIM providers.

1. Enter the endpoint for managing users in the **Users Resource URL** field.
2. Enter the endpoint for managing groups in the **Group Resource URL** field if the SCIM service provider supports this notion and you want to provision groups as well.
3. Select the authentication method the endpoint accepts.

Select **None** if the endpoint does not require authentication.

4. If you select either **Basic Authentication** or **OAuth 2.0 Bearer Token** as the authentication method, enter valid credentials in the **User** and **Password** fields for your provider.

Additionally, if you choose **OAuth 2.0 Bearer Token** as the authentication method, in order to support the password grant type, enter valid credentials in the **Client ID** and **Client Secret** fields and the corresponding URL in the **Token Endpoint URL** field.

5. If the SCIM provider does *not* support PATCH updates, clear the associated check box.

For information about PATCH, see the [SCIM specification](http://www.simplecloud.info/specs/draft-scim-api-01.html#edit-resource-with-patch) (www.simplecloud.info/specs/draft-scim-api-01.html#edit-resource-with-patch).


6. Optional: Depending on the target-provider needs, select the option to provision groups by supplying complete LDAP DNs, rather than only common names (CNs), to identify groups.

Some SCIM providers, including PingOne, allow administrators to parse full DN's when necessary (for example, in the case of duplicate CN's) to determine group access mapping to specific applications based on other DN elements.


7. Select a deprovision method (**Delete User** or **Disable User**).

Deprovisioning is triggered when previously provisioned users no longer meet the condition set in the **Source Location** screen.

- The **Delete User** option removes the user account.
- The **Disable User** option deactivates the user account.


 **Note:** For SaaS provisioning, the provisioner does *not* necessarily delete deprovisioned users from target data stores in accordance with common practice. Rather their status is changed to indicate that the accounts are no longer active.

8. Click **Next**.

 **Note:** For some provisioning plug-ins (including the built-in SCIM outbound provisioner), when you first enter or change credentials and click **Next**, PingFederate immediately tests connectivity to the target.

Specify custom SCIM attributes

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** screen.

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 Multi-Valued Attributes
- Complex Attributes
- Complex Multi-Valued 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 Multi-Valued	Not Applicable
options	Complex	quantity, strike, first, and last
tablets	Complex Multi-Valued	model, serial, and type.

 **Note:** type is a reserved sub-attribute for a complex multi-valued attribute.

 **Tip:** For more information about SCIM and attribute types, see the website www.simplecloud.info.


1. Specify the URI of the schema extension in the **Extension Namespace** field.



 **Tip:** 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.

 **Tip:** 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:

Task

Change the attribute name

Steps

1. Replace the current value in the **Name** field.
2. Click **Done**.



Set the attribute as a simple multi-valued attribute

1. Select the **Is Multi-Valued** check box.
2. Click **Done**.

Task

Steps

SP Connection | Configure Channels | Custom SCIM Attribute

Custom SCIM Attribute Options

Custom SCIM Attribute Configuration. You are presented with options to make the attribute complex and/or multi-valued.

NAME:

Sub-Attributes

Sub-Attributes	Action
<input type="text"/>	<input type="button" value="Add"/>

IS MULTI-VALUED

Types

Types	Action
<input type="text"/>	<input type="button" value="Add"/>

Add sub-attributes to make the attribute a complex attribute

1. Enter a sub-attribute and click **Add**. (Repeat this step to add more sub-attributes as needed.)
2. Use the **Edit**, **Update**, and **Cancel** workflow to make or undo a change to the name of a sub-attribute. Use the **Delete** and **Undelete** workflow to remove a sub-attribute or cancel the removal request.
3. Click **Done**.

SP Connection | Configure Channels | Custom SCIM Attribute

Custom SCIM Attribute Options

Custom SCIM Attribute Configuration. You are presented with options to make the attribute complex and/or multi-valued.

NAME:

Sub-Attributes

Sub-Attributes	Action
price	<input type="button" value="Edit"/> <input type="button" value="Delete"/>
rate	<input type="button" value="Edit"/> <input type="button" value="Delete"/>
tax	<input type="button" value="Edit"/> <input type="button" value="Delete"/>
<input type="text"/>	<input type="button" value="Add"/>

IS MULTI-VALUED

Add sub-attributes and set the attribute as a complex multi-valued attribute

1. Enter a sub-attribute and click **Add**. (Repeat this step to add more sub-attributes as needed.)



Tip: Use the **Edit**, **Update**, and **Cancel** workflow to make or undo a change to the name of a sub-attribute. Use the **Delete** and **Undelete** workflow to remove a sub-attribute or cancel the removal request.

2. Select the **Is Multi-Valued** check box.
3. Specify at least one value under the Types column for `type`, a reserved sub-attribute for a complex multi-valued attribute.



Tip: Use the **Edit**, **Update**, and **Cancel** workflow to make or undo a change to the `type` value. Use the **Delete** and **Undelete** workflow to remove a `type` value or cancel the removal request.

4. Click **Done**.

SP Connection | Configure Channels | Custom SCIM Attribute

Custom SCIM Attribute Options

Custom SCIM Attribute Configuration. You are presented with options to make the attribute complex and/or multi-valued.

NAME:

Sub-Attributes

Sub-Attributes	Action
model	<input type="button" value="Edit"/> <input type="button" value="Delete"/>
rate	<input type="button" value="Edit"/> <input type="button" value="Delete"/>
tax	<input type="button" value="Edit"/> <input type="button" value="Delete"/>
<input type="text"/>	<input type="button" value="Add"/>

IS MULTI-VALUED

Types

Types	Action
type	<input type="button" value="Edit"/> <input type="button" value="Delete"/>
<input type="text"/>	<input type="button" value="Add"/>

When you finish editing custom attributes, click **Next** in the **Custom SCIM Attributes** screen.

Manage channels

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; for example, if your organization has separate LDAP stores (or different nodes in the same store) for various user groups needing SSO access and provisioning to the same domain.



Tip: There can be only one target domain for provisioning per connection. If your organization subscribes to multiple domains for which you need provisioning support, you need a separate SP connection for each domain.




On the **Manage Channels** screen, you can perform the following tasks:

- Click **Create** to add a new channel.
Alternatively, you create a new channel by copying an existing channel (and making other required changes).
- Select an existing channel to modify its settings.
- Use the **Delete** and **Undelete** workflow to remove or cancel the removal request of existing channel.

Specify channel information

In the **Channel Info** screen, 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.



1. Enter a channel name.

If you are copying a channel, you must enter a new value in the **Channel Name** 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**.

Identify the source data store

PingFederate[®] supports PingDirectory[™], Microsoft Active Directory, and the 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 `<pf_install>/pingfederate/server/default/conf/template/ldap-templates` directory.)

Information from your user-data store is used to supply mapped values for each user attribute required by the SP.



1. On the **Source** screen, choose the LDAP store to use for this channel.

If the data store 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 data store.

2. Click **Next**.

Modify source settings

The **Source Settings** screen shows the default configuration of the data store selected on the **Source** screen, 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 data store representing the user's Globally Unique Identifier.
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 multi-value user attribute containing the DNs of the groups to which an entry belongs. This attribute does not apply to some LDAP servers, including the Oracle Directory Server. The attribute below is used instead. Active Directory uses both values to provide a two-way mapping between User and Group objects.
Group Member Attribute	The name of a multi-value 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: <p>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.</p> <p>Timestamp – Queries for timestamps on user records that are <i>not older</i> than the last time records were checked. This check is more efficient from the point of view of the PingFederate provisioner but can be more time consuming on the LDAP side, particularly with the Oracle Directory Server.</p> <p>Timestamp No Negation – Queries for timestamps on user records that are <i>newer</i> than the last time records were checked. This algorithm is recommended for the Oracle Directory Server.</p>
USN Attribute	The name of the attribute used to store the update sequence number—applicable when the Active Directory algorithm is chosen above.
Timestamp Attribute	The name of the attribute used to store the timestamp on user records.

Field	Description
Account Status Attribute	The name of the attribute in which the user's account status (active or inactive) is stored, for example, Active Directory = <code>userAccountControl</code> and Oracle Directory Server = <code>nsaccountlock</code> .
Account Status Algorithm	The method by which PingFederate determines a user's account status. The values are: Active Directory Bitmap – For Active Directory, which uses a bitmap for each user entry. For more information about <code>userAccountControl</code> flags, see Microsoft's knowledge base (support.microsoft.com/kb/305144). Flag – For Oracle Directory Server and other LDAP directories that use a separate attribute to store the user's status. When this option is selected, the Flag Comparison Value and Flag Comparison Status fields below are also used.
Default Status	Indicates the user's status if the attribute is missing.
Flag Comparison Value	Indicates the value for the attribute (for example, <code>nsaccountlock</code>) that PingFederate expects to be returned. Used when the Account Status Algorithm is set to Flag .
Flag Comparison Status	Indicates whether the user is enabled or disabled when the flag has the value specified in the Flag Comparison Value field. Setting the value to <i>true</i> equals enabled while setting the value to <i>false</i> equals disabled. For example, if the Account Status Attribute is set to <code>nsaccountlock</code> , and the Flag Comparison Value is set to true , and the Flag Comparison Status is set to false , then any users with <code>nsaccountlock=true</code> are disabled. Used when the Account Status Algorithm is set to Flag .

If you are using PingDirectory™, Microsoft Active Directory, and the Oracle Directory Server, in most cases no changes are needed on this screen unless your data store uses a customized schema.

If you are using a different LDAP directory, you must supply the required information on this screen unless you have defined a template for the data store. (For more information, see the `sample.template.txt` in the `<pf_install>/pingfederate/server/default/conf/template/ldap-templates` directory.)

1. Modify the settings, as needed.
2. Click **Next**.

Specify a source location

Indicate on the **Source Location** screen where PingFederate® should look for user records in the data store. The same location may be used to retrieve user-group DNs for maintaining corresponding groups at the service provider.

The screenshot shows a web interface for configuring the source location. It has a breadcrumb trail: "SP Connection | Configure Channels | Channel". Below that are tabs: "Channel Info", "Source", "Source Settings", "Source Location", "Attribute Mapping", and "Attribute & Security". The "Source Location" tab is active. The main content area has a heading: "Enter the Base DN where user records are located in the data store, and specify where to LDAP Filter or Group DN." There are two sections: "Users" and "Groups". Each section has a "BASE DN" field, a "GROUP DN" field, and a "FILTER" field. There are also checkboxes for "NESTED SEARCH" in each section. At the bottom, there are buttons for "Cancel", "Save Config", "Previous", and "Next".

After specifying the required base DN, you have the options to provision users (and groups when applicable) based on group membership information or LDAP search results.









Note: Groups provisioning is supported for SCIM and the Google Apps Connector (version 2.0 and higher) but may not be supported for other SaaS Connectors. If not, the associated fields under **Groups** on the **Source Location** screen are inactive. Support for the feature may become available in future SaaS Connector releases; please refer to documentation in your add-on distribution package.

1. Enter the base DN where user records are stored in the **Base DN**.

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, as described in the following table:


Object	Field description
Users	<p>Group DN</p> <p>The distinguished name (DN) of a group in the user repository whose <i>member users</i> should be provisioned.</p> <p> Important: If a Group DN is specified in conjunction with a Microsoft Active Directory data store, the domain functional level must be set to Windows Server 2008 R2 or Windows Server 2012.</p> <p>Refer to Windows Server documentation or support for more information.</p> <p>(Optional) Select the Nested Search check box to include <i>users</i> 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).</p> <p> Note: The Nested Search feature is available when Microsoft Active Directory or Oracle Directory Server is selected as the source user repository (see Identify the source data store on page 320).</p> <p>Filter</p> <p>An LDAP search filter that returns <i>user objects</i> representing the <i>users</i> that should be provisioned.</p> <p>For information about LDAP filters, refer to your LDAP documentation. Note that you may need to escape any special characters.</p> <p> Important: The Group DN field is ignored when a Filter field value is configured.</p>
Groups (when applicable)	<p>Group DN</p> <p>The distinguished name (DN) of the group in the user repository that should be provisioned.</p> <p> Important: If a Group DN is specified in conjunction with a Microsoft Active Directory data store, the domain functional level must be set to Windows Server 2008 R2 or Windows Server 2012.</p> <p>Refer to Windows Server documentation or support for more information.</p> <p>(Optional) Select the Nested Search check box to include <i>groups</i> 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).</p> <p> Note: The Nested Search feature is available when Microsoft Active Directory or Oracle Directory Server is selected as the source user repository (see Identify the source data store on page 320).</p> <p>Filter</p> <p>An LDAP search filter that returns <i>group objects</i> representing the <i>groups</i> that should be provisioned.</p>

Object	Field description
	For information about LDAP filters, refer to your LDAP documentation. Note that you may need to escape any special characters.
	 Important: The Group DN field is ignored when a Filter field value is configured.


3. Click **Next**.

Map attributes


The **Attribute Mapping** screen provides a means of managing how attributes from your user store are mapped to SCIM attributes in the core schema and custom attributes through a schema extension or to the provisioning fields supported for your organization's SaaS-customer account.




Field Name	Attribute(s)	Expression	Default	Action
Username*	uid		uid	Edit
Password*	password		password	Edit
Entry Name	cn		cn	Edit
Given Name	givenName		givenName	Edit

 **Important:** If you are provisioning for SCIM, your SP may make one or more optional core attributes mandatory. Refer to the SCIM documentation from the SP or the SCIM Resource Schema representation.

For SaaS Connectors, refer to the connection-template setup steps in the *Quick Connection Guide* of the SaaS Connector for information about any special fields and how to map them.


 **Tip:** For non-SCIM SaaS connectors, PingFederate® automatically retrieves from the vendor the **Field Names** shown on this screen, but only on the first pass through the screen flow. If you are using this screen to modify an existing mapping configuration, click **Refresh Fields** near the bottom of the screen to synchronize the list with the target if needed.

For each field, the **Specify Attribute Mapping** screen 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** near the of the screen 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.

1. Click **Edit** under Action for a field.

 **Tip:** If you have specified any custom attributes, they are listed at the end of the **Attribute Mapping** screen.

2. On the **Specify Attribute Mapping** screen, provide mapping details.

3. Repeat for each attribute shown in the **Field Name** column as needed.


 **Tip:** For most fields, if you need to map more than one attribute from your data store into a single field at the target location, then you must use an OGNL expression to indicate how the attribute values are to be combined.

The only exception is a field called LDAP Attributes Map, provided primarily to support PingOne®-specific SCIM attributes. This field may contain multiple attributes without using OGNL.

4. Click **Next**.

Specify mapping details

On the **Specify Attribute Mapping** screen, you define specific mapping information for each field required for provisioning (and for any optional fields, as needed).

 **Caution:** 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.

Define mapping information for a standard attribute

1. Optional: 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 may 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](#). You can also do both for certain attributes, as needed: that is, specify both LDAP attributes and a default value.

2. Select the source attribute from the class under **Attribute** and then click **Add Attribute**.
3. Repeat the steps above to add additional applicable attributes, as needed, to use in a mapping expression.



Important: You must add an attribute for it to be used in an expression.

4. Enter or select a default value under **Value Definition** (optional, if one or more attributes is specified above).

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 evaluating the expression.

5. If more than one attribute is used for mapping fields other than LDAP Attributes Map, enter an expression.



Tip: Click **Edit** to create and validate the expression.

6. Optional: Select one or more processing options, as defined below:

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 value(s).

Mask Log Values – Determines whether sensitive information (for example, the Password attribute) will be masked in PingFederate® log files, or not.

Upper Case, Lower Case, or None – Transforms the attribute value(s) 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 (for example, 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.




Tip: For SaaS Connectors not bundled with PingFederate, refer to your SaaS Connector *Quick Connection Guide* for instructions on mapping options or requirements for particular provisioning fields.


7. Click **Done**.

Define mapping information for a custom attribute:

1. Select a sub attribute under **Attribute ID**.

 **Note:** Applicable only to *Complex Attributes* or *Complex Multi-Valued Attributes* (see [Specify custom SCIM attributes](#) on page 317).


2. Optional: 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 may 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](#). You can also do both for certain attributes, as needed: that is, specify both LDAP attributes and a default value.

3. Select the source attribute from the class under **LDAP Attribute** and then click **Add Attribute**.

4. Optional: Select one or more processing options, as defined below:

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 value(s).

Mask Log Values – Determines whether sensitive information (for example, the Password attribute) will be masked in PingFederate® log files, or not.

Upper Case, Lower Case, or None – Transforms the attribute value(s) 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 (for example, 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.

 **Tip:** For SaaS Connectors not bundled with PingFederate, refer to your SaaS Connector *Quick Connection Guide* for instructions on mapping options or requirements for particular provisioning fields.

5. Optional: Enter a default value.

6. Click **Add Mapping**.

 **Note:** For complex attributes or complex multi-valued attributes, repeat these steps to map additional sub attributes as needed.


7. Click **Done**.

Review channel settings

When you finish setting up a channel, you may choose to activate it immediately; or you can return to the **Activation & Summary** screen and activate the channel when needed. Note that the SP connection must also 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 SSO and SLO transactions may still occur (if an associated connection is active).

- To toggle the status, select **Active** or **Inactive**, and then click **Save**.

 **Caution:** When a channel is activated, initial provisioning occurs as soon as the synchronization-frequency time period expires (see [Configure outbound provisioning settings](#) on page 155). The default

is 60 seconds. Initial provisioning can consume considerable processing time, depending on the amount of data that needs to be transmitted; administrators may wish to plan accordingly.

- To modify channel settings, click the associated heading.



Important: Regardless of whether you choose to activate a new channel immediately or later, if you want to save the channel configuration, click **Save** in the **Activation & Summary** screen.

Review SP connection settings

When you finish creating or modifying a connection, you can review the connection settings and toggle the connection status on the **Activation & Summary** screen.



Important: When creating a new connection, the default connection status is **Inactive** when you reach the **Activation & Summary** screen.

Regardless of whether you choose to activate a new connection now or later, you must click **Save** on the **Activation & Summary** screen if you want to keep the new connection.

The SSO application endpoint (underneath the **Connection Status** field) is a sample URL that webmasters or web application developers at your site might use to invoke SSO for the connection (see [View IdP application endpoints](#) on page 280).

- If you need to make any changes, click the heading over the information you want to edit.
- If you are creating a new connection and want to keep your configuration, you must click **Save**.
- If you are modifying an existing connection (including toggling the connection status) and want to keep your changes, you must click **Save**.

Define SP affiliations

An 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 (see [Account linking](#) on page 69).

This may be of use when multiple service providers 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 IdP adapter instance for generating account links (see [Manage authentication source mappings](#) on page 294).

You can create or modify an SP affiliation from the Main Menu under IdP Configuration or from a list of affiliations (see [Manage SP affiliations](#) on page 327).

To create an SP affiliation:

1. Click **IdP Configuration** on the Main Menu.
2. Click **Create New** under SP Affiliations.



Tip: See [Import affiliation metadata](#) on page 328 for subsequent steps.

To modify or delete an affiliation:

- See [Manage SP affiliations](#) on page 327.

Manage SP affiliations

You can manage SP affiliations on this screen.

To reach this screen for editing:

1. Click **IdP Configuration** on the Main Menu.
2. Click **Manage All** under SP Affiliations.

To begin creating a new affiliation:

- Click **Create Affiliation** (see the next sections for more information).

To delete an affiliation:

1. Click **Delete** under Action for the affiliation you want to delete.
2. Click **Save** to confirm the deletion (or click **undelete**).

To view or modify an affiliation:

- Click the affiliation ID.

Import affiliation metadata

An IdP may send a metadata file containing information that automatically specifies members of an SP affiliation for use in PingFederate.

- If you do not have a metadata file, click **Next**.

To import metadata:

1. Click **Browse** to locate and import the file and then click **Next**.
2. Review the information on the Create Affiliation page (see the next section).
3. Click **Save** on the Summary screen.

Enter affiliation information

Enter or modify basic information about an affiliation on the Affiliation General Info screen.

If you imported a metadata file, this information is already supplied. However, you may change the Affiliation ID or select a different Affiliation Owner, if required.

Field 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 may serve as the Owner.

Manage affiliation membership

On the Affiliation Membership screen, you 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 may add or remove connections from the affiliation.


- To add an SP partner connection to the affiliation, select the connection from the drop-down list and click **Add**.
 - ⚠ **Important:** Each connection in the affiliation must be configured to use the same IdP adapter instance for generating account links (see [Manage authentication source mappings](#) on page 294).
- To remove a member of the affiliation, click **Delete** under Action for the connection and click **Save**.
 - 📄 **Note:** 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.

Review an SP affiliation

From the Affiliation Management Summary screen you can activate or deactivate an SP affiliation. You also save new affiliations on this screen, or you can click heading links to go back and modify information.

To change an Affiliation Status:

- Select either Active or Inactive and then click **Save**.

 **Important:** Be sure to click **Save**. Otherwise, the status will not be changed.

To edit a connection:

1. Click the heading above the information you want to modify.
2. Make your change and click **Save**.

Configure SP Auto-Connect

When your SP partner is also using PingFederate 5 or higher (or is otherwise able to provide interoperable SAML 2.0 metadata via HTTP on demand), you may choose to use Auto-Connect for that partner (see [Using Auto-Connect](#)). This configuration can be shared among an unlimited number of SAML 2.0 partners.



Note: You enable the SAML 2.0 Auto-Connect profile under System Settings (see [Choose roles and protocols](#) on page 151).

Once Auto-Connect is enabled on your PingFederate server, complete the configuration from the Main Menu under IdP Configuration. This configuration entails:

- Setting up a common connection for all Auto-Connect partners
- Establishing a list of SP partner domains authorized to use the connection

Initial setup for SP Auto-Connect

The basic configuration for Auto-Connect requires only:

- Defining a period of validity for assertions (assertion lifetime)
- Choosing a signing certificate for assertions and other SAML messages
- Configuring assertion-creation information

All other partner-connection specifications are handled automatically at runtime.

Specify an assertion lifetime for SP Auto-Connect

Identity-federation standards require a window of time during which an SSO token is considered valid. Each SSO token has an issuance time-stamp element and elements indicating the allowable lifetime of the SSO token (in minutes) before and after the issuance time stamp.

Field Descriptions

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.

To change the default times:

- Optional: Edit the desired setting(s) and click **Next** or **Save**.

Choose a signing certificate for SP Auto-Connect

For Auto-Connect runtime processing, assertions and SLO messages must be signed, because they are sent over either the POST or redirect bindings (see [SAML 2.0 profiles](#) on page 34 in the “Supported Standards” chapter of the *Get started with PingFederate* guide).



Note: The signing certificate is embedded in your server's Auto-Connect metadata (see [Using Auto-Connect](#) on page 79); there is no need to exchange certificates with your partners.

You can use the same certificate used for signing metadata (see [Configure metadata signing](#) on page 156). If you use a different certificate, ensure that it meets Auto-Connect validation requirements (see [Auto-Connect security model](#) on page 80).

To specify a certificate:

1. Select a signing certificate from the list.

If you have not yet created or imported your certificate into PingFederate, click **Manage Certificates** (see [Manage digital signing certificates and decryption keys](#) on page 184).

2. 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.

Configure assertion creation for SP Auto-Connect

Configuring assertion creation for Auto-Connect is similar to configuring the same settings for regular partner connections.

- Click **Configure Assertion Creation** to continue.

For configuration information, refer to sections under [Configure SSO token creation](#) on page 291.

Review SP Auto-Connect settings

When you finish configuring your SP Auto-Connect initial setup, you may choose to activate the common connection immediately on the Activation & Summary screen. (No runtime processing occurs until your partner's Auto-Connect gateway is also established and a user initiates an SSO or SLO event.)

- ⚠ **Important:** Regardless of whether you choose to activate a newly configured connection now or later, you must click **Save** on the Activation & Summary screen if you want to keep the configuration.

You can deactivate the connection at any time (for maintenance, for example). While a connection is inactive, all SSO or SLO transactions to or from Auto-Connect partners are disabled.

To change a Connection Status:

- Select Active or Inactive and then click **Save**.

To modify a setting:

1. Locate the currently configured setting under one of the summary headings and then click the subheading above the data.

📄 **Note:** Changes made to Auto-Connect settings will be out of sync, temporarily, with metadata caches that any currently active partners might be using. If your connection is in production, you might wish to lower your server's metadata lifetime in advance of making configuration changes (see [Configure metadata lifetime](#) on page 157).

2. Change the information and click **Save**, if available.

If **Save** is not available, additional, dependent changes are required; click **Next** or **Done** until you reach a screen containing a **Save** button. Then click **Save** and continue as needed until you return to the Main Menu.

Specify allowed SP domains

This screen provides PingFederate® with a list of trusted domain names of your Auto-Connect partners.

Normally, when PingFederate receives an authentication request from a domain in this list), the runtime engine completes the connection automatically using metadata obtained from a standard, public location— `http://saml.<domain_name>`. (See [Using Auto-Connect](#).) Alternatively, if an Auto-Connect partner elects not to use the standard location, you can supply the applicable URL.

Service provider SSO configuration

In an SP role, you use the PingFederate® administrative console to configure local application-integration information and to manage connections to your IdP-partner sites. Prior to configuring connections to IdPs, you must establish your site as an SP on the **Server Configuration > Server Settings > Roles & Protocols** screen.

Note that only one connection is needed per partner, even if you are integrating more than one web application.

While your entity ID is defined on the **Server Configuration > Server Settings > Federation Info** screen, you may identify your organization differently through the use of virtual server IDs on a per-connection basis (see [Multiple virtual server IDs](#) on page 84).

Additionally, you may deploy an SP connection to bridge a service provider to one or more identity providers through one or more authentication policy contracts (see [Federation hub use cases](#) on page 86 and [Federation hub and authentication policy contracts](#) on page 89 for more information).



Note: This topic applies to configuration settings needed for browser-based SSO. While there is some cross-over information also applicable to WS-Trust STS, if you are using PingFederate *exclusively* as an STS, start with [WS-Trust STS configuration](#) on page 400.

SP application integration settings

The integration of local applications with PingFederate® is the essential "last-mile" configuration that allows end-users at your IdP partner's website to access your protected resources. This process is facilitated through the use of application-integration kits and a robust Software Development Kit (see [SSO integration kits and adapters](#) on page 67).

Under **Application Integration** on the **SP Configuration** menu, you configure the SP adapters that PingFederate uses to create user sessions that allow SSO access to your protected resources. You can also set Default URLs to which users may be directed during SSO or SLO, and you can 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 STS, the selections under **Application Integration** also include menu items for configuring plug-in **Token Generators** (see [SP configuration for STS](#) on page 412).

Manage SP adapters

An SP adapter is used to create a local-application session for a user in order for PingFederate® to provide SSO access to your applications or other protected resources. You must configure at least one instance of an SP adapter in order to set up connections to IdP partners. You can also 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 also deploy additional integration kits from the Ping Identity [Downloads](#) website.

You manage SP adapter instances on the **SP Configuration > Adapters** screen.

- 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**.



Important: After installing new adapter program files, you may be required to make additional configuration changes in areas such as adapter instances and connections, as prompted by the administrative console.



Note: Automatic multi-connection error checking occurs by default whenever you access this screen. The intent is to verify that configured connections have not been adversely affected by changes made here.

If you experience noticeable delays in accessing this screen, you can optionally disable automatic connection validation on the **Server Configuration > Server Settings > System Options** screen.

Create an SP adapter instance


The first step in creating an adapter instance is choosing an adapter type.

- On the **Type** screen, configure the basics of this adapter instance.
 - a) Enter the required information and select the adapter type from the 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 screens.

Configure an SP adapter instance

Depending on the selected adapter, the **Instance Configuration** screens presents you with different configuration parameters.

- Follow the on-screen instructions to configure the adapter instance.
 -  **Note:** If this is a child instance, select the override check box to modify the configuration.
- For information about configuring the OpenToken Adapter distributed with PingFederate, see [Configure the OpenToken SP Adapter](#) on page 444.
- If you have deployed an adapter from an integration kit, locate the user guide from our [documentation website](#) and configure the adapter instance accordingly.


Invoke SP adapter actions

Adapters can be written to 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.

- Follow the on-screen instructions to complete the actions required.

Extend an SP adapter contract

Adapters may be written with an option allowing administrators to add to the attributes required for creating usable sessions. This feature might be needed, for example, 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.

- Enter the name of the desired attribute and click **Add**.

Repeat this step as needed to add another attribute.

Review an SP adapter configuration

1. On the **Summary** screen, review your configuration, modify as needed, and click **Done** to exit the **Create Adapter Instance** workflow.
2. On the **Manage SP Adapter Instances** screen, click **Save** to retain the configuration of the adapter instance.


If you want to exit without saving the configuration, click **Cancel**.

Configure target URL mapping

When you have more than one target session defined in an IdP connection, you must map the target URL to its target session. When PingFederate® receives an SSO or 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.


For example, this mapping configuration may 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. It is worth noting that when this information is provided with the SP request, the mapping table is ignored (see [SP services](#) on page 451).

Furthermore, 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.

 **Tip:** In this scenario, PingFederate is a federation hub for the identity provider and the service providers (see [Federation hub use cases](#) on page 86).

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 **SP Configuration > Target URL Mapping** screen. The configuration process involves entering an URL and select a target session for it. Refer to the following table for more information.

Field	Description
URL	<p>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.</p> <p>You may use a wildcard (*) to match multiple URLs to the same target session but you can use only one wildcard (*) per URL.</p> <p>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.</p> <p> Note: If a target session is not allowed based on restrictions imposed (see Restrict a target session to certain virtual server IDs on page 350), PingFederate tries the next entry.</p>
Target Type	The type of the Target Session . 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 .
Target Session	A selection of configured SP adapter instances or SP connections. The available values depends on the chosen the Target Type list.

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:// sp.example.com/ acct101/	OpenToken SP Adapter to an local training app
http:// sp.example.com/*	SP connection to SP SaaS

A target URL of `http://sp.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://sp.example.com/*`) matches the target URL.

1. Enter a URL.
2. Select a target type from the list.

Applicable only when the IdP role is activated with at least one protocol for browser-based SSO.

3. Select a target session from the list.
4. Click **Add Mapping**.
5. Repeat these steps to add multiple mappings.

Use the up and down arrows to re-arrange the order of the mappings. Use the **Edit**, **Update**, and **Cancel** workflow to make or undo a change to a mapping. Use the **Delete** and **Undelete** workflow to remove a mapping or cancel the removal request.

Configure Identity Store Provisioners

PingFederate allows you to create custom Identity Store Provisioners to bridge the inbound 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 application-specific user database schema.

Using the Software Developer Kit for PingFederate, you can create and test these custom Identity Store Provisioners (see the PingFederate [SDK developer's guide on page 559](#)).

To support custom attributes, you must add the schema extension and the custom attributes to the 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 [Define custom SCIM attributes](#) on page 374 and [SCIM 1.1 Core Schema](#) (www.simplecloud.info/specs/draft-scim-core-schema-01.html).



Note: The Identity Store Provisioner option is active only after you enable the Inbound Provisioning protocol on the Roles & Protocol screen (see [Choosing Roles and Protocols](#)).



Note: Automatic multi-connection error checking occurs by default whenever you access this screen. The intent is to verify that configured connections have not been adversely affected by changes made here.

If you experience noticeable delays in accessing this screen, you can optionally disable automatic connection validation on the **Server Configuration > Server Settings > System Options** screen.

Create an Identity Store Provisioner instance

On the Type screen, you begin creating an instance of an identity store that PingFederate uses to bridge the inbound SCIM processing to an external user store using a custom implementation.

Field descriptions

Field	Description
Instance Name	A descriptive name for the provisioner instance—for example, an identity management system name.
Instance ID	An internal identifier for the provisioner instance. Must be alphanumeric with no spaces.
Type	A list of deployed provisioner types available for creating a provisioner instance for the server. A developer typically deploys any new provisioner types before an administrator sets up a connection partner. The Identity Store Provisioner Type is limited to the provisioners currently installed on your server.
Parent Instance (Optional)	A list of configured instances for this type of provisioner. If applicable, select an instance that can be used as a basis for a parent/child configuration.

To define an identity store provisioner:

1. Click **SP Configuration** on the Main Menu.
2. Click **Identity Store Provisioners** under Application Integration Settings.
3. Click **Create New Instance on the Manage Identity Store Provisioners screen**.
4. Enter the **Instance Name** and **Instance Id** on the Type screen.
5. Select the **Type** from the drop-down menu.
6. Optional: Select a **Parent Instance** from the 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 overridden. You can specify overrides during the rest of the setup.


7. Click **Next** and enter information on subsequent screens for this identity store setup, as indicated in the following sections.
8. Click **Done** on the Summary screen.
9. Click **Save** on the **Manage Identity Store Provisioners** screen.

Define the Identity Store Provisioner behavior

Different configuration parameters are available on the Identity Store Provisioner screen. These options are controlled by the provisioner plug-in software (see topics about identity store provisioner interfaces in the PingFederate *SDK developer's guide* for more information).

Extend the Identity Store Provisioner contract


Identity Store Provisioners may be written with an option allowing administrators to add to the core attributes the plug-in instance requires. Both the core and extended contract attributes you define here must be mapped when you configure “Write Users” within an Inbound Identity Store Provisioner connection.

-  **Tip:** To keep your plug-in 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.

To reach this screen for editing:

1. Click **SP Configuration** on the Main Menu.
2. Click **Identity Store Provisioners** under Application Integration Settings.
3. Click an Instance Name.
4. Click **Extended Contract** (if available).


To add an attribute:

-  **Note:** If this is a child instance, select the override check box to modify the configuration.

1. Enter the attribute name in the text box and click **Add**.
2. Click **Done** then click **Save** on the Manage Identity Store Provisioners screen.

Extend the Identity Store Provisioner contract for groups


Identity Store Provisioners may be written with an option allowing administrators to add to the core group attributes the plug-in instance requires. Both the core and extended group attributes you define here must be mapped when you configure “Write Groups” within an Inbound Identity Store Provisioner connection.

-  **Tip:** To keep your plug-in 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.

To reach this screen for editing:

1. Click **SP Configuration** on the Main Menu.
2. Click **Identity Store Provisioners** under Application Integration Settings.
3. Click an Instance Name.
4. Click **Extended Group Contract** (if available).

To add an attribute:

-  **Note:** If this is a child instance, select the override check box to modify the configuration.

Enter the attribute name in the text box and click **Add**.

Click **Done** then click **Save** on the Manage Identity Store Provisioners screen.

Review the Identity Store Provisioner configuration

From the Summary screen, you can reach identity store provisioner instance settings for editing.

To edit the configuration:

1. Click the heading above the information you want to change.
2. Make your changes.
3. Click **Done** on the configuration page and **Save** on the Manage Identity Store Provisioners screen.

To save an identity store provisioner instance:


1. Click **Done** on the Summary screen.
2. Click **Save** on the Manage Identity Store Provisioners screen.

Configure default URLs


As an SP, you can supply a default URL that the end-user may see when an SSO request succeeds (that is, a session is created at your site) but the target resource is not available or not specified.

 **Note:** You can also specify default target SSO URLs for individual IdP connections, which take precedence over this global setting (see [Configure default target URLs](#) on page 363).

Similarly, you can specify to prompt a default URL indicating a successful SLO to the end-user (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 [Localize messages for end users](#) on page 139. If localization is not needed, you may also specify a message directly in this field to change the default.

Your application or your partner's application may supply these URLs at runtime (see [SP services](#) on page 451), but if none is provided, PingFederate will use the default values you enter on this screen unless, in the case of SSO, a default is also defined for the connection.

 **Tip:** If no default targets are specified here or at the connection level (for SSO), PingFederate provides built-in landing pages for the user. These web pages are is among the templates you can modify with your own branding or other information (see [Customizable user-facing screens](#) on page 126).

View SP application endpoints

Web-application developers at your site need to know the application endpoints to initiate transactions via PingFederate.

- Click **Application Endpoints** on the **SP Configuration** menu 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 [SP services](#) on page 451 and [System-services endpoints](#) on page 460.


Federation settings

If your identity federation uses the SAML 2.0 XASP profile (see [Attribute Query and XASP](#) in the “Supported Standards” chapter of the *Get started with PingFederate* guide), you may need to identify the IdP connection to which an attribute request applies. If so, click **Attribute Requester Mapping** under Federation Settings on the SP Configuration menu to complete the configuration (see [Manage attribute requester mappings](#) on page 337).

Also under Federation Settings, you can view protocol endpoints that your federation partners need to know to access your services via PingFederate.

Manage 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 (see [Attribute Query and XASP](#) on page 43). Optionally, an application may also supply an issuer DN, which can be used to determine the correct IdP (*Attribute Authority*) to use for a set of users associated with an IdP.

 **Note:** The Format query parameter must be set to a specified value for XASP (see [SP services](#) on page 451).

You can map X.509 identifying information to connections and specify a default connection on the **SP Configuration > Attribute Requester Mapping** screen. (Note that this screen is only presented if the XASP profile is enabled on the **Server Configuration > Server Settings > Roles & Protocols** screen.)

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 query 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 may use a regular expression to match different DNs to the same connection. Only one expression may be used in any single entry. DN values must be entered in all lower-case characters.

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 an 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.

Use the **Edit**, **Update**, and **Cancel** workflow to make or undo a change to an entry. Use the **Delete** and **Undelete** workflow to remove an entry or cancel the removal request.

View SP protocol endpoints

Click **Protocol Endpoints** in the **SP Configuration** menu 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 **Server Configuration > Server Settings > Federation Info** screen. These endpoints are built into PingFederate and cannot be changed.


Your federation partners or STS clients need to know the applicable SP services endpoints to communicate with your PingFederate server. Configured service endpoints for SAML connections are included in metadata export files.


PingFederate provides a favorite icon for all protocol endpoints. For more information, see [Customize the favicon for application and protocol endpoints](#) on page 144.

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:

Service	URL and Description
	<ul style="list-style-type: none"> 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 login_hint parameter with a default value to the IdP connection on the OpenID Provider Info screen. 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 login_hint parameter in the URL. <p>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 (see Configure request parameters and SSO URLs on page 395).</p> <p>For more information about Third Party Initiated Login flow, see the OpenID Connect specification (openid.net/specs/openid-connect-core-1_0.html#ThirdPartyInitiatedLogin).</p>
Single Logout Service (SAML 2.0)	<p>/sp/SLO.saml2</p> <p>The URL that receives and processes logout requests and responses.</p>
Assertion Consumer Service (SAML 2.0)	<p>/sp/ACS.saml2</p> <p>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.</p>
Artifact Resolution Service (SAML 2.0)	<p>/sp/ARS.ssaml2</p> <p>The SOAP endpoint that processes artifacts returned from a federation partner to retrieve the referenced XML message on the back channel. (See Important note at the end of this table.)</p>
Metadata Service	<p>/</p> <p>The default endpoint (empty path) from which partners can retrieve Auto-Connect metadata.</p>
Assertion Consumer Service (SAML 1.x)	<p>/sp/acs.saml1</p> <p>A SAML 1.x implementation URL that receives and processes assertions from an IdP.</p>
Single Sign-on Service (WS-Federation)	<p>/sp/prp.wsf</p> <p>The WS-Federation implementation URL that receives and processes security tokens and SLO messages.</p>
WS-Trust STS (two endpoints)	<p>/sp/sts.wst</p> <p>The SOAP endpoint that receives and processes SAML security-token requests from STS clients (Web Service Providers at the SP site), validating SAML tokens or validating and exchanging SAML tokens based on the configured IdP connection.</p> <p>/pf/sts.wst</p>

Service	URL and Description
	<p>Initiates direct STS token-to-token exchange and token validation, from an IdP token processor to an SP token generator, when that feature is configured (see Token translator mappings on page 423).</p> <p> Note: If multiple token-generator instances of the same type are configured for the same connection or token-to-token mapping, a query parameter, <code>TokenGeneratorId</code>, must be added to either of these endpoints—see Manage token generators on page 412.</p> <p>(See Important note at the end of this table.)</p>

 **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` (see [Configure PingFederate properties](#) on page 93).

Virtual server ID support

For SAML connections using multiple virtual server IDs (see [Multiple virtual server IDs](#) on page 84), each virtual server ID has its own set of protocol endpoints. You may export a connection metadata for your partner on the **Server Configuration > Metadata Export** screen (see [Provide SAML metadata by file](#) on page 116).

For WS-Federation (and SAML) connections using multiple virtual server IDs, you may provide your partner the federation metadata endpoint (`/pf/federation_metadata.ping`) with the `PartnerIdpId` and `vsid` parameters; for example:

Partner's entity ID	Your virtual server ID	Federation metadata URL
SP	idev1	<code>https://www.example.com/pf/sts_mex.ping?PartnerIdpId=IdP&vsid=idev1</code>
	idev2	<code>https://www.example.com/pf/sts_mex.ping?PartnerIdpId=IdP&vsid=idev2</code>

(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 may 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 460.)

Note that the virtual server ID concept does not apply to the `/pf/sts.wst` endpoint because token-to-token exchange does not involve any connections. As needed, you may pass the token-to-token endpoint to your partners as-is.

Manage IdP connections

As an SP, you manage connection settings to support the exchange of federation-protocol messages (OpenID Connect, SAML, WS-Federation, or WS-Trust) with an IdP, OAuth client, OpenID Provider (OP), or STS client application at your site.

These settings include:

- User attributes that you expect to receive in an SSO token (SAML assertion or WS-Trust STS SAML token).
- User attributes that 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 (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 (see [Federation planning checklist](#) on page 82).

As an SP, you respond to user requests for SSO and 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 (see [Manage target session mappings](#) on page 348).

(No adapter instance or authentication policy contract is required for a connection that uses only the Attribute Query profile. For more information, see [Manage Attribute Query profile](#) on page 365.)

If you intend to pass attribute values to an adapter instance from a local data store, you must define the data store during this configuration, if you have not done so already (see [Manage data stores](#) on page 158).

Administrative interface

You manage connection settings using the **IdP Connection** wizard, 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 screen, where you manage one or more settings. You may move to a sibling task using the **Next** or **Previous** button. If you are on a sub task, you may also move to its parent task using the **Done** button.

When creating a new connection, you may save your progress using the **Save Draft** button. Note that not all screens offer this option. When you reach the **Activation & Summary** screen, you must click **Save** to complete the new connection.

When editing an existing connection, you may make changes and then click **Save** to commit your changes. In order words, you are not required to step through all screen to reach the **Activation & Summary** screen before you can save your changes.



Note: The **Save** button is available on most screen. If a screen does not show a **Save** button, click **Next** or **Done** until you reach to a screen where you can use its **Save** button to commit your changes.

Access IdP connections

The **SP Configuration** menu displays a list of the most-recently modified IdP connections. You may create or import a connection. You may also edit a recently modified connection by clicking on its connection name.

To access the rest of the connections, click **Manage All** to open the **IdP Connections** screen.

IdP Connections

The **IdP Connections** screen displays 20 connections at a time. As needed, you can use the pagination controls to navigate through the rest of your connections. You can also search connections by their names or connection IDs.



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.

You can sort by connection name, partner's connection ID, and default virtual server ID; narrow by protocol and status; and perform the following actions.

Create a connection

Follow the **Connection** wizard to create a new connection to your SP partner.

Copy a connection

Follow the **Connection** wizard to create a new connection based on an existing (source) connection. This is most useful if the new connection and the source connection have many setting values in common. Note that the new connection stays inactive until you change its status.

Export a connection

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

Follow the **Import Connection** wizard to import a connection export. If the connection already exists, you have the option to overwrite the existing connection.



Note: Prior to the import, you may modify the XML file to suit your needs. The XML file may also be imported to another PingFederate® environment acting in the same federation role (IdP or SP) at your site. Note that the source and the target must run the same version of PingFederate.

Export SAML metadata

Follow the **Export Metadata** wizard to generate a metadata for any SAML Browser SSO connection.

Update a SAML connection with metadata

Follow the **Import Metadata** wizard to update a SAML Browser SSO connection with metadata. You may update a connection via a metadata XML file or a metadata URL.



Important: The update operation may require additional configuration. Reviewing the connection after the update operation is recommended.

Delete a connection

Delete an inactive connection. You must click **Save** on the **SP Connections** screen to confirm your request. If you change your mind, click **Undelete** to cancel your request.

Override logging mode

Override the verbosity of runtime transaction logging for all SP connections.

Edit a connection

Open the connection by clicking on its name, select the setting that you want to reconfigure, and complete the change.

Resolve IdP connection errors

PingFederate® automatically validates configured connections before displaying them on the **Connections** screen. 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.

- To resolve the error, select the connection and follow the on-screen instructions to modify the configuration, one connection at a time.

Choose an IdP connection type

Indicate on the **Connection Type** screen whether the connection to this partner is for Browser SSO, WS-Trust 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.

If you have selected multiple protocols on the **Server Configuration > Server Settings > Roles & Protocols** screen, you must select the applicable protocol on the **Connection Type** screen 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. Note that each connection must use a unique (partner) connection ID.

- To configure a connection for secure browser-based SSO, select the **Browser SSO Profiles** check box and a protocol from the list (if necessary).
- To configure an STS connection, select the **WS-Trust STS** check box and the default token type from the list.
- To configure a connection that exchanges SAML assertions or JWTs for access tokens, select the **OAuth Assertion Grant** check box.

Note that the **OAuth Assertion Grant** option is available only if at least one Access Token Manager instance has been configured on the **OAuth Settings > Access Token Management** screen.

For more information about these standards, see [Security Assertion Markup Language \(SAML\) 2.0 Profile for OAuth 2.0 Client Authentication and Authorization Grants](https://tools.ietf.org/html/rfc7522) (tools.ietf.org/html/rfc7522) and [JSON Web Token \(JWT\) Profile for OAuth 2.0 Client Authentication and Authorization Grants](https://tools.ietf.org/html/rfc7523) (tools.ietf.org/html/rfc7523).

- To configure an inbound provisioning connection, make that selection 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) are preserved.

Note that the **Inbound Provisioning** option is active only if the **Inbound Provisioning** role is enabled on the **Server Configuration > Server Settings > Roles & Protocols** screen.

- Optional: If your PingFederate® license manages connections by groups, then you can select a group for this connection.

This option is not displayed for unrestricted or other types of licenses.

Choose IdP connection options

On the **Connection Options** screen (shown only for browser-based SSO connections), you can enable browser-based SSO in conjunction with JIT provisioning. Additionally, you may also choose to map user attributes for persistent grants used by the optional PingFederate® OAuth authorization server.

For SAML 2.0, you also have the option of configuring the **Attribute Query** profile, with or without the browser-based SSO.

- To create a connection for browser-based SSO, select the **Browser SSO** check box.
- To enable JIT provisioning, OAuth attribute mapping, or both for this connection, make that selections (after selecting the **Browser SSO** check box).


Note that the **OAuth Attribute Mapping** option is only available when the OAuth 2.0 Authorization Server (AS) role is enabled on the **Server Configuration > Server Settings > Roles & Protocols** screen.

- To create a connection to facilitate the SAML 2.0 Attribute Query profile, select the **Attribute Query** check box (see [Attribute Query and XASP](#) on page 43).

Import IdP metadata

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
- Enter a new metadata URL
- Select an existing metadata URL

 **Tip:** Using a metadata URL streamlines the configuration process. For example, by importing a metadata URL, you can quickly establish a Browser SSO connection to an InCommon-participating partner (see www.incommon.org/participants).


When you enter a new URL or select an existing metadata URL, PingFederate® also enables the automatic update option and checks the metadata daily. If PingFederate detects changes in the partner's digital signing certificates, encryption key, or contact information, it updates the connection automatically. If you prefer to update the connection manually, you can clear the **Enable Automatic Reloading** check box.

The frequency is configurable through the **Reload Delay** field on the **Server Configuration > Server Settings > Metadata Lifetime** screen.






Although optional, it is recommended that you turn on email notifications for SAML metadata update events on the **Server Configuration > Server Settings > Runtime Notification** screen.



 **Note:** If the metadata contains changes that require additional configuration, the email message also provides a list of the applicable items.

After the connection is created, you can add, remove, or change the metadata URL associated with the connection in the **Import Metadata** screen. In addition, you can also toggle the **Enable Automatic Reloading** option for the connection.

 **Tip:** These capabilities also lower the cost of maintaining connections with InCommon participants.

Refer to the following steps to import or update metadata. Instructions vary depending on the medium of the metadata.

Metadata medium	Steps
A metadata file	<ol style="list-style-type: none"> On the Import Metadata screen, select the File option. Choose the metadata file, and then click Next. <ul style="list-style-type: none">  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 screen, and then click Next. On the Metadata Summary screen, review the signature information to evaluate the authenticity of the metadata. Click Next.
A new metadata URL	<ol style="list-style-type: none"> On the Import Metadata screen, select the URL option. Enter a new metadata URL. <ul style="list-style-type: none">  Note: If you specify a URL that has already been entered into the system through the Server Configuration > Metadata URLs workflow, the administrative console reminds you to select it from the Existing URL Name list. Optionally, clear the Enable Automatic Reloading check box to disable automatic update. Click Load Metadata. <ul style="list-style-type: none">  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 screen, and then click Next. On the Metadata Summary screen, review the signature information to evaluate the authenticity of the metadata. Click Next.
An existing metadata URL	<ol style="list-style-type: none"> On the Import Metadata screen, select the URL option. Select an existing metadata from the list, and then verify the URL in the Selected Existing URL. Optionally, clear the Enable Automatic Reloading check box to disable automatic update. Click Load Metadata.

Metadata medium	Steps
	<p> Note: If the metadata contains multiple entries, select the desired partner from the Select Entity ID list and click Next.</p> <p> Note: If there is a digital signature error, correct the issue using the Server Configuration > Metadata URLs workflow.</p> <p>5. Click Next.</p>

Identify the partner

On the **General Info** screen, 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 screen you can define a default error message that end users see in the event that SSO fails.

1. Provide the basic information to identify your partner.

Refer to the following table for more information.

Field	Description
Partner's Entity ID, Issuer, Partner's Realm, or Connection ID (Required)	<p>The published, protocol-dependent, unique identifier of your partner.</p> <p>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 may have been obtained out-of-band or via a SAML metadata file.</p> <p>For a WS-Federation connection, this is your partner's Realm.</p> <p>For an OpenID Connect connection, this is the Issuer Identifier of the OpenID Provider (OP).</p> <p>For an STS-only connection, this ID can be any unique identifier.</p>
Connection Name (Required)	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.
Virtual Server IDs (Not applicable to OpenID Connect connections)	<p>If you want to identify your server to this connection partner using an ID other than the one you specified on the Server Configuration > Server Settings > Federation Info screen, enter a virtual server ID in this field and click Add.</p> <p>Enter additional virtual server IDs as needed.</p>
Client ID and Client Secret (Applicable to and required for OpenID Connect connections)	<p>The client ID and the client secret to communicate with the OP.</p> <p>This client represents PingFederate and is created and managed at the OP. For more information, please refer to the documentation provided by the OP.</p>
Base URL	The fully qualified hostname and port on which your partner's federation deployment runs (for example, <code>https://www.example.com:9031</code>). 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.

Field	Description
Error Message	<p>If an error occurs on this server, the end user's browser may be redirected (in a default situation) to an error page hosted within PingFederate.</p> <p>The default entry is a variable found in the <pf_install>/pingfederate/server/default/conf/language-packs/pingfederate-messages.properties file, which is part of the PingFederate localization framework. If localization is not needed, you may also specify a message directly in this field to change the default.</p> <p>(Applicable only to IdP connections using the SAML or OpenID Connect browser SSO profile.)</p>
Logging Mode	The level of transaction logging applicable for this connection.

If the OP supports the OpenID Connect Discovery specification, the connection setup can be 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.

In addition, the protocol settings can be refreshed 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 could be mapped to the target applications later. In the rare 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 may 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**.

(Note that this step is not applicable to an STS-only connection.)

Configure Browser SSO

Browser-based SSO (also known as Browser SSO) is another term for secure SSO, which relies on a user's web browser and HTTP to broker SSO tokens (in contrast to WS-Trust STS messaging, which is typically application-driven across the back channel and does not require browser mediation).

- To continue, click **Configure Browser SSO**.

SP Browser SSO configuration steps

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 (see [Federation roles](#) on page 30). 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 SSO steps

- [Select SAML profiles](#) on page 347
- [Configure user-session creation](#) on page 347
 - [Choose an identity mapping method](#) on page 347
 - [Define an attribute contract](#) on page 348
 - [Manage target session mappings](#) on page 348
- [Manage protocol settings](#) on page 357
 - [Specify SSO service URLs \(SAML\)](#) on page 357

- [Define SLO service URLs \(SAML 2.0\)](#) on page 359
- [Select allowable SAML bindings \(SAML\)](#) on page 359
- [Specify an artifact lifetime \(SAML 2.0\)](#) on page 360
- [Define artifact resolver locations \(SAML\)](#) on page 360
- [Configure default target URLs](#) on page 363
- [Override authentication context in an IdP connection](#) on page 363
- [Configure signature policy](#) on page 364
- [Specify XML encryption policy \(for SAML 2.0\)](#) on page 364

WS-Federation SSO steps

- [Configure user-session creation](#) on page 347
 - [Choose an identity mapping method](#) on page 347
 - [Define an attribute contract](#) on page 348
 - [Manage target session mappings](#) on page 348
- [Manage protocol settings](#) on page 357
 - [Specify a service URL \(WS-Federation\)](#) on page 358
 - [Configure default target URLs](#) on page 363
 - [Configure signature policy](#) on page 364

SAML 1.x SSO steps

- [Select SAML profiles](#) on page 347
- [Configure user-session creation](#) on page 347
 - [Choose an identity mapping method](#) on page 347
 - [Define an attribute contract](#) on page 348
 - [Manage target session mappings](#) on page 348
- [Manage protocol settings](#) on page 357
 - [Specify SSO service URLs \(SAML\)](#) on page 357
 - [Select allowable SAML bindings \(SAML\)](#) on page 359
 - [Define artifact resolver locations \(SAML\)](#) on page 360
 - [Configure default target URLs](#) on page 363
 - [Configure signature policy](#) on page 364

OpenID Connect SSO steps

- [Configure user-session creation](#) on page 347
 - [Choose an identity mapping method](#) on page 347
 - [Define an attribute contract](#) on page 348
 - [Manage target session mappings](#) on page 348
- [Manage protocol settings](#) on page 357
 - [Configure OpenID Provider information](#) on page 360
 - [Configure default target URLs](#) on page 363
 - [Override authentication context in an IdP connection](#) on page 363

After configuring SSO settings, you will normally need to configure authentication credentials, the range of which depends on your SSO selections (see [Configure security credentials](#) on page 385). Also, other configuration tasks may remain to be configured for new or modified connections, depending on the selected options on the **Connection Options** screen.

Select SAML profiles

A SAML profile is the message-interchange scenario that you and your federation partner have agreed to use.

For SAML 2.0, PingFederate® supports all IdP- and SP-initiated SSO and SLO profiles. For SAML 1.x, PingFederate supports both the standard IdP-initiated SSO profile and a proprietary ("destination-first") SP-initiated SSO profile.

(For information on typical SAML SSO and SAML 2.0 SLO profile configurations, including illustrations, see [SAML 1.x profiles](#) on page 32 and [SAML 2.0 profiles](#) on page 34.)

Note that the **SAML Profiles** screen does not apply to OpenID Connect and WS-Federation IdP connections.

- Select the applicable profile (or profiles) based on your partner agreement.

For SAML 2.0, you must always 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 screen is **SP-initiated SSO**.

Configure user-session creation

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 may associate one or more authentication policy contracts to the IdP connection (see [Federation hub use cases](#) on page 86 for more information).

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.

- To continue, click **Configure User-Session Creation**.

Choose an identity mapping method

PingFederate® allows an SP to use either account linking or account mapping to associate remote users with local accounts for SSO between business partners (see [Identity mapping](#) on page 69). On the **Identity Mapping** screen, you choose which method to use in this IdP connection. You and your partner should decide in advance which option to use (see [Federation planning checklist](#) on page 82).

If your site is using account linking, then establishing an attribute contract is not required. Depending on your partner agreement, however, you may 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 (see [About attributes](#) on page 70).

- ⚠ **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.

- If you want to dynamically associate remote users with local accounts using a known attribute to identify a user (for example, 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.

- 📘 **Tip:** If you are using PingFederate's JIT provisioning, choose **Account Mapping** (see [Configure just-in-time provisioning](#) on page 366).

- 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.

- 📘 **Tip:** PingFederate uses a default, HyperSQL database to handle account linking. You can use your own database instead, as needed. For more information, see [Define an account-linking data store](#) on page 166.

- 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**. Additionally, select **No Mapping** if you are deploying an IdP connection solely for OAuth attribute mapping without the use of an authentication policy contract (see [Configure OAuth attribute mapping](#) on page 354).


Define an attribute contract

An attribute contract is the set of user attributes that you and your partner have agreed will be sent in SSO tokens for this connection (see [Attribute contracts](#) on page 71). You may extend the attribute contract with additional attributes. Optionally, you can configure PingFederate® to mask individual extended attributes in its logs (see [Attribute masking](#) on page 74).

-  **Tip:** If you are creating or updating a SAML or an OpenID Connect 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.

1. 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 (see www.incommon.org/participants), the assertion may 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](http://tools.ietf.org/html/rfc4524) (tools.ietf.org/html/rfc4524) and [RFC4519](http://tools.ietf.org/html/rfc4519) (tools.ietf.org/html/rfc4519). The following table describes a subset of the OIDs (object IDs) 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	o
2.5.4.42	givenName
2.16.840.1.113730.3.1.241	displayName

For other attributes, refer to 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.

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.


Manage target session mappings

Remote users arriving at your site via an SSO request do so in order to use specific applications or gain access to protected resources. Based on the nature of the business relationship and the agreement with your partner, you may be


expected to provide access to these applications. Therefore, integration between your federation SP server and local applications is important. PingFederate® uses SP adapters to identify the user to your application based on attributes sent in an SSO token. A configured adapter is known as an adapter instance.

In a basic scenario, you map an SP adapter instance to an IdP connection on the **Target Session Mapping** screen and complete its mapping configuration through a series of sub tasks. 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.

 **Tip:** To learn more about authentication policies, see [Authentication policies](#) on page 209.

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 which you map values from the SSO tokens.


 **Tip:** To learn more about federation hub, see [Federation hub use cases](#) on page 86.

On the **Target Session Mapping** screen (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), you should map multiple SP adapter instances. If you are bridging an identity provider to multiple service providers, you should map multiple authentication policy contracts.

Note that the **Target Session Mapping** configuration does not apply when the **No Mapping** option is chosen on the **Identity Mapping** screen.

- To map an SP adapter instance, click **Map New Adapter Instance**.
- To map an APC, click **Map New Authentication Policy**.
- To 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.
- To 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**.
- If you are editing an existing configuration and want to keep your changes, click **Save**.

When target sessions are restricted to certain virtual server IDs, the allowed IDs are displayed under **Virtual Server IDs**.

 **Note:** 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 (see [Configure target URL mapping](#) on page 332).

Select a target session


The first task of the mapping configuration is to map an adapter instance or an authentication policy contract to your connection.

- To map an adapter instance, follow these steps:
 - a) Select an adapter instance from the list.

If you do not see the desired adapter instance, click **Manage Adapter Instances** to create a new instance of any deployed adapter.

- b) Select the **Override Instance Settings** check box if you want to customize one or more adapter settings for this connection alone.

When selected, the administrative console adds a new set of sub tasks (the **Override Instance** screen and its sub tasks).

 **Tip:** 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 plug-in configurations](#) on page 69.

- To map an APC, select an adapter instance from the list.


If you do not see the desired APC, click **Manage Authentication Policy Contracts** to create a new policy contract.

If you are editing a currently mapped adapter 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 adapter settings specifically for this connection.

If you are editing a currently mapped APC, no changes can be made on this screen.


Override an SP adapter instance

On the **Override Instance** screen, 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.

- Click **Override Instance Settings**.

On each of the settings screens, select the **Override** check box, make your changes, and then click **Next**. When you are finished, click **Done** to continue with the rest of the mapping configuration.

 **Note:** The override setting screens are functionally identical to those used for creating a new adapter instance (see [Manage SP adapters](#) on page 331).

If you are editing a currently mapped adapter instance, click **Override Instance Settings** to reconfigure any overridden settings for this connection. You may also remove all overridden settings on a per-screen basis by clearing the **Override** check box near the top of the screen.

Restrict a target session to certain virtual server IDs

When you multiplex one connection for multiple environments (see [Multiple virtual server IDs](#) on page 84), you can enforce integration requirements by restricting a target session to certain virtual server IDs on the **Virtual Server IDs** screen. By default, no restriction is imposed.

1. Select the **Restrict Virtual Server IDs** check box.
2. Select one or more virtual server IDs that you want to allow for this target session.


If you are editing a currently mapped adapter instance or APC, you can toggle the **Restrict Virtual Server IDs** setting. You can also change the allowed virtual server IDs.

Choose an attribute mapping method

On the **Adapter Data Store** screen for SP adapter mapping (or the **Attribute Retrieval** screen for authentication policy contract mapping), you select if and how PingFederate® should query a local data store to help fulfill the attribute contract in conjunction with attribute values from the SSO token.

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 data store can supply it.

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 configure data store queries as part of the fulfillment configuration for the applicable APC.

 **Tip:** To learn more about authentication policies, see [Authentication policies](#) on page 209.

- If the attribute contract contains all the attributes that your application requires, select **Use only the attributes available in the SSO assertion**.

- To set up a data store query, select **Use the SSO assertion (or provider claims) to look up additional information** and then follow a series of sub tasks to complete the configuration.

For step-by-step instructions, see [Choose a data store](#) on page 508.

If you are editing a currently mapped adapter instance or APC, you can change the mapping method, which may require additional configuration changes in subsequent tasks.

Configure target session fulfillment

On the **Adapter Contract Fulfillment** screen, 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.

If you are bridging an identity provider to one or more service providers, the values mapped to the authentication policy contracts are used by the associated SP connections to create assertions for the service providers (see [Federation hub use cases](#) on page 86).

At runtime, an SSO operation fails if PingFederate® cannot fulfill the required attribute.

- For each attribute, select a source from the list and then choose or enter a value.

- **Account Link**

When selected, the **Value** list is populated 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. In addition, this source appears only if you have elected to use account linking for a target session on the **Identity Mapping** screen.

- **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 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 is populated with the available context of the transaction. Select the desired context from the list.



Note: The **HTTP Request** context value is retrieved as a Java object rather than text. For this reason, OGNL expressions are more appropriate to evaluate and return values (see **Expression**).

- **LDAP, JDBC, or Custom**

When selected, the **Value** list is populated with attributes that you have selected from the data store. Select the desired attribute from the list.

- **Expression** (when enabled)

This option provides more complex mapping capabilities; for example, 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 (see **Text**).

Note that expressions are not enabled by default. For more information about enabling and editing OGNL expressions, see [Attribute mapping expressions](#) on page 497.

- **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 data store, when applicable, using this syntax:

```
${ds.attribute}
```

where *attribute* is any attribute that you have selected from the data store.



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).

There are a variety of reasons why you might hard code a text value. For example, if your web application provides a consumer service, you might want to supply a particular promotion code for the partner.


If you are editing a currently mapped adapter instance or APC, you can update the mapping configuration, which may require additional configuration changes in subsequent tasks.

Identify issuance criteria

On the **Issuance Criteria** screen, you may define the criteria that PingFederate® can evaluate to determine whether the user is authorized to continue the SSO transaction using the SP adapter or authentication policy contract. This optional token authorization configuration can be used to restrict who can SSO to protected resources. (For more information, see [About token authorization](#) on page 74.)

Issuance criteria and OGNL expressions


When defining multiple criteria using the user-interface elements, all criteria must pass in order for a user to be authorized. Source attribute values are evaluated against specified values using an exact match algorithm. The **multi-value contains** (or **multi-value does not contain**) conditions (intended for attributes that may contain multiple values) return true if the source attribute contains the exact specified value. For more complex evaluations, including conditional criteria and partial matching, define criteria using OGNL expressions.

 **Important:** When you multiplex one connection for multiple environments (see [Multiple virtual server IDs](#) on page 84), consider using an OGNL expression to verify the virtual server ID in conjunction with other conditions, such as group membership information, to protect against unauthorized access (see [Sample OGNL expressions](#) on page 498).

Note that expressions are not enabled by default. For more information about enabling and editing OGNL expressions, see [Attribute mapping expressions](#) on page 497.

1. Select the source of the attribute that you want PingFederate to evaluate from the **Source** list.

When you select a source, the **Attribute Name** list is populated with the associated attributes or properties, as described in the following table.


Source	Attribute name
Assertion or Provider Claims	Populated with the attributes from the SSO token.
Context	Populated with the available context of the transaction.  Note: The HTTP Request context value is retrieved as a Java object rather than text. For this reason, OGNL expressions are more appropriate to evaluate and return values.
LDAP, JDBC, or Custom (if a data store is configured)	Populated with attributes that you have selected from the data store.
Mapped Attributes	Populated with attributes from the attribute contract.

2. Select an attribute from the **Attribute Name** list.
3. Select the evaluation method from the **Condition** 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

 **Note:** Use one of the *multi-value* conditions when you want PingFederate to verify that an attribute contains or does not contain the specified value in its attribute list. Using a single-value condition when an attribute has multiple values causes the condition to evaluate consistently to false.

4. Enter the exact value in the **Value** field.

Wildcards (for partial matching) are not allowed.

 **Tip:** You can use OGNL expressions to perform more complex evaluations, including partial matching.

5. Optional: Use the **Error Result** field to customize an error code or error message for use if the token authorization fails.


Error results are handled in one of two ways:

Redirect

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 `sp.sso.error.page.template.html` template file.

 **Tip:** You may also enter an error code in the **Error Result** field that the error template or the web application can use to process the error result in a variety of ways; for example, a redirect to another page or e-mail to an administrator.

If you leave this field blank, a default `ACCESS_DENIED` error result is used if the authorization fails.


6. Click **Add**.

7. Repeat these steps if you want to add additional criteria.


 **Important:** All criteria must pass in order for a user to be authorized.

8. Optional: Click **Show Advanced Criteria** to configure more complex evaluations using OGNL expressions.

- Enter an OGNL expression in the **Expression** field.
- Enter an error code or an error message in the **Error Result** field.

 **Note:** If the OGNL expression resolves to a string value (rather than `true` or `false`), the returned value overrides the **Error Result** field value.

- Click **Add**.
- Click **Test**, enter values in the applicable fields, and test the resulting output for the expression.
- Repeat these steps if you want to add additional criteria using expressions.

 **Important:** All criteria must pass in order for a user to be authorized.

If you want conditional criteria, add only one criterion with such processing logic.

If you are editing a currently mapped adapter instance or APC, you can add, remove, or edit any issuance criteria.

Review the target session mapping

When you have finished adding a new or modifying an existing instance of an adapter or an authentication policy contract, you can review the configuration on its **Summary** screen.

- If you need to make any changes, click the heading over the information you want to edit.

- If you are creating a new connection and want to keep your configuration, click **Done**.
- If you are editing an existing configuration and want to keep your changes, click **Save**.


Review the session creation summary

When you are finished with the **User-Session Creation** task, you can review the configuration on its **Summary** screen.

- If you need to make any changes, click the heading over the information you want to edit.
- If you are creating a new connection and want to keep your configuration, click **Done**.
- If you are editing an existing configuration and want to keep your changes, click **Save**.

Configure OAuth attribute mapping

This configuration allows administrators to use assertion attributes from an IdP connection for mapping values into OAuth persistent-grant USER_KEY and extended attributes.

1. Create a new IdP connection or select an existing IdP connection from the **SP Configuration** menu.
2. On the **Connection Type** screen, select the **Browser SSO Profiles** check box and the applicable protocol.
3. On the **Connection Options** screen, select the **Browser SSO** check box and then select the **OAuth Attribute Mapping** check box.
 -  **Tip:** You may also select other options on the **Connection Type** and **Connection Options** screens. 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** screen, enter the required information.
5. On the **Browser SSO** screen, click **Configure Browser SSO** and follow the wizard to complete the **User-Session Creation** configuration.
6. On the **OAuth Attribute Mapping** screen, 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) on the **User-Session Creation > Target Session Mapping** screen, you may select the **Map to OAuth via Authentication Policy Contract** option, and then select the applicable APC from the list.

Choose an OAuth data store

This optional configuration is the same for all OAuth attribute-mapping task flows.


- If you do not need additional attributes from a data store, just click **Next** on the Data Store screen.

To reach this screen for editing

1. Click **SP Configuration** on the Main Menu.
2. Click the connection name under IdP Connections.
 - Click **Manage All**, if needed, to see a full list of connections.
3. Click **Browser SSO** under the IdP Connection tab.
4. Click **Configure Browser SSO**.
5. Click **OAuth Attribute Mapping** under the Browser SSO tab.
6. Click **Configure OAuth Attribute Mapping**.
7. Click **Data Store** on the Summary screen.

Configure contract fulfillment

The last step in configuring OAuth attribute mapping is to map attributes for the USER_KEY (to look up persistent OAuth grants), the extended attributes (to fulfill the access token), and USER_NAME (the name shown to end-users on permissions screens) for persistent access-token grants.

-  **Note:** The `USER_KEY` values must be unique across all end users because it is the identifier to store and to retrieve persistent grants. If `SAML_SUBJECT` uniquely identifies all end users, you can map `SAML_SUBJECT` to `USER_KEY`.


Map each attribute from one of the following Sources:

- Assertion

Values are contained in the assertions from this IdP. When you make this selection, the associated Value drop-down list is populated by the attribute contract.

- **Context**

When selected, the **Value** list is populated with the available context of the transaction. Select the desired context from the list.

-  **Note:** The **HTTP Request** context value is retrieved as a Java object rather than text. For this reason, OGNL expressions are more appropriate to evaluate and return values (see **Expression**).

- JDBC/LDAP/Custom

Values are returned from your user-data store. When you make this selection, the Value list is populated by the LDAP, JDBC or Custom attributes identified for this data store.

- **Expression** (when enabled)

This option provides more complex mapping capabilities; for example, 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 (see **Text**).

Note that expressions are not enabled by default. For more information about enabling and editing OGNL expressions, see *Attribute mapping expressions* on page 497.

- 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 data store, when applicable, using this syntax:

```
${ds.attribute}
```

where `attribute` is any of the data store attributes you have selected.

1. Choose a source and then choose (or enter) a value for each attribute in the contract.
2. Click **Next**.

To reach this screen

1. Click **SP Configuration** on the Main Menu.
2. Click the connection name under IdP Connections.
Click **Manage All**, if needed, to see a full list of connections.
3. Click **Browser SSO** under the IdP Connection tab.
4. Click **Configure Browser SSO**.
5. Click **OAuth Attribute Mapping** under the Browser SSO tab.
6. Click **Configure OAuth Attribute Mapping**.
7. Click **Contract Fulfillment** on the Summary screen.


Specify issuance criteria for OAuth attribute mapping

Use the **Issuance Criteria** screen to define criteria PingFederate® can evaluate to determine whether to issue an access token for a user. This token authorization can be used to restrict who can access a resource.


To configure Issuance Criteria:

1. Select the source of the attribute that you want PingFederate to evaluate from the **Source** list.

When you select a source, the **Attribute Name** list is populated with the associated attributes or properties, as described in the following table.


Source	Attribute name
Assertion or Provider Claims	Populated with the attributes from the SSO token.
Context	Populated with the available context of the transaction.  Note: The HTTP Request context value is retrieved as a Java object rather than text. For this reason, OGNL expressions are more appropriate to evaluate and return values.
LDAP, JDBC, or Custom (if a data store is configured)	Populated with attributes that you have selected from the data store.
Mapped Attributes	Populated with attributes from the attribute contract.

2. Select an attribute name.
3. Select the condition you want to apply from the list.

 **Note:** Use one of the *multi-value* conditions when you want PingFederate to verify that an attribute contains or does not contain the specified value in its attribute list. Using a single-value condition when an attribute has multiple values causes the condition to evaluate consistently to false.

4. Enter the exact value in the **Value** field.

Wildcards (for partial matching) are not allowed.


 **Tip:** You can use OGNL expressions to perform more complex evaluations, including partial matching.

5. Optional: Use the **Error Result** field to customize an error code or error message for use if the token authorization fails.


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.

If you leave this field blank, the user sees a default ACCESS_DENIED error result at runtime if the authorization fails.


6. Click **Add**.
7. Repeat the steps above to add additional attributes, as needed, for each authorization criteria.

 **Note:** All criteria must pass in order for a user to be authorized.

8. Optional: Click **Show Advanced Criteria** to configure more complex evaluations using OGNL expressions.

 **Note:** Expressions must be enabled for the **Show Advanced Criteria** button to appear.

- Use the in-line editor to enter the OGNL expression.
- Use the **Error Result** field to enter an error message or an error code for use if authorization fails (see step 5 for more information).

 **Note:** If the OGNL expression resolves to a string value (rather than true or false), then the returned value overrides the Error Result.

- Click **Add**. Continue to add expressions, as needed.
- Click the **Test** link to input values and test the resulting output for the expression.

9. Click **Next** (or **Done**, if you are finished with the configuration, and then **Save** on the **Manage Mappings** screen).

To edit Issuance Criteria:

- Click **Edit** under Actions for the criteria.

To delete Issuance Criteria:

Click **Delete** under Actions for the criteria and then click **Save**.

Review the OAuth attribute mapping summary

When you finish the configuration, you can review it on the **Summary** screen.

If you need to make any changes, click the heading over the information you want to edit. When you finish, click **Done**.

Manage protocol settings

The **Protocol Settings** screen provides the launching point for configuring partner endpoints, message customizations, and other protocol-specific settings for browser-based SSO connections.

SAML 2.0

- Outbound SSO bindings (POST, redirect, artifact) and the corresponding SSO service URLs
- Outbound 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 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 OP in its authorization and token requests
- The OpenID Connect login type and authentication scheme used by PingFederate when communicating with the OpenID Provider
- The authorization endpoint, the token endpoint, the user information (UserInfo) endpoint, and the JWKS URL
- Default target URL
- Authentication context mappings
- To continue, click **Configure Protocol Settings**.

Specify SSO service URLs (SAML)

The SSO service endpoint is a location to which PingFederate® sends authentication requests when SSO is initiated at your site, according to partner requirements. It is applicable to all SAML versions when the SP-initiated SSO profile is enabled.

For SAML 2.0 connections, you associate bindings to the endpoints where your 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 may 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.

- a) Enter the SSO service endpoint to the **Endpoint URL** field.

You may enter a relative path (begin with a forward slash) if you have provided a base URL on the **General Info** screen.

For SAML 1.x connections, this is the only configurable item on the **SSO Service URL** screen.

The remaining steps on the **SSO Service URLs** screen are only applicable 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.

2. Optional: Customize messages using OGNL expressions.

Note that expressions are not enabled by default. For more information about enabling and editing OGNL expressions, see [Attribute mapping expressions](#) on page 497. In addition, message customization is not applicable to SAML 1.x connection.

- a) Click **Show Advanced Customizations**.
 b) 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 [Customize assertions and authentication requests](#) on page 501.

- d) Click **Add**.
 e) Optional: Repeat to add another message customization.

Specify a service URL (WS-Federation)

The service endpoint URL is a location to which PingFederate® sends RST (Request for Security Token) and SLO messages.

To protect against session token hijacking, PingFederate provides an option to validate wreply for SLO. When the option is enabled, you can specify additional allowed domains and paths in this screen. 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.



Note: Settings to enter additional allowed domains and paths appear only if the option to validate wreply for SLO is enabled (see [Manage partner redirect validation](#) on page 177).

1. Enter the WS-Federation protocol endpoint at the IdP site in the **Endpoint URL** field.

You may enter a relative path (begin with a forward slash) if you have provided a base URL on the **General Info** screen.

2. Specify additional allowed domains and paths.

- a) Optional: Indicate whether to 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.

- b) Enter the expected domain or IP address under **Valid Domain Name**.


Use the domain only, without qualifiers; for example: `example.com`

* `.example.com` covers `hr.example.com` or `email.example.com` but not `example.com`.

 **Important:** While using an initial wildcard provides the convenience of allowing multiple sub domains with one entry, consider adding individual sub domains to limit the redirection to a list of known hosts.


- c) Optional: Enter the exact (and case-sensitive) path of the resource under **Valid Path**.

Starts 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. For example, `/app/Consumer.jsp` allows `/app/Consumer.jsp` but rejects `/app/consumer.jsp`.

 **Tip:** You can also enter multiple query parameters with or without a fragment.

For instance, `/app/Consumer.jsp?area=West&team=IT#ref1001` matches `/app/Consumer.jsp?area=West&team=IT#ref1001` but not `/app/Consumer.jsp?area=East&team=IT#ref1001`.

- d) Optional: Select the check box under **Allow Any Query/Fragment** if you want to allow any query parameters or fragment in the resource.

 **Note:** When selected, no query parameter, fragment, or a combination of them is allowed in the path.

- e) Click **Add**.

- f) Optional: Repeat the previous step to add additional entries as needed.

Note that a more specific match is considered a better match and an exact match is considered the best match.

Define SLO service URLs (SAML 2.0)

On the **SLO Service URLs** screen, you associate bindings to the endpoints where your IdP receives logout requests when SLO is initiated at your site and where PingFederate® sends SLO responses when it receives SLO requests from the IdP.

This step applies only to SAML 2.0 connections when either SLO profile is selected on the **SAML Profiles** screen.

1. Select a SAML binding from the list; for example, **POST**.
2. Enter the SLO endpoint URL to the **Endpoint URL** field.

You may enter a relative path (begin with a forward slash) if you have provided a base URL on the **General Info** screen.

3. Optional: Enter an URL to the **Response URL** field.

When specified, it 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 may enter a relative path (begin with a forward slash) if you have provided a base URL on the **General Info** screen.

4. Click **Add**.
5. Optional: Repeat to add additional SLO endpoints.

If you are editing an existing connection, you can reconfigure the SLO endpoints, which may require additional configuration changes in subsequent tasks.

Select allowable SAML bindings (SAML)

On the **Allowable SAML Bindings** screen, you select the one or more bindings that your IdP partner can use to send SAML assertions or SAML 2.0 SLO messages.

This configuration applies to all SAML connections.

- Select only the applicable SAML binding (or bindings) based on your partner agreement.

If you have specified an SSO or SLO endpoint using the artifact (outbound) binding, you must including SOAP as one of the allowable (inbound) binding.

If you are editing an existing connection, you can reconfigure the allowable bindings, which may require additional configuration changes in subsequent tasks.

Specify an artifact lifetime (SAML 2.0)

When PingFederate® sends an artifact to your IdP's SSO or SLO service endpoint, an element in the message indicates how long it should be considered valid. On the **Artifact Lifetime** screen, specify the expiry information in seconds.

You can change the default value per your requirements, if needed. Also consider synchronizing clocks between your server and your partner's SAML gateway server. If clocks are not synchronized, you might need to set the artifact lifetime to a higher value.

This step applies only to SAML 2.0 connections.

- Optional: Override the default value of the **Artifact Lifetime** field.

The default value is 60 (seconds).

If you are editing an existing connection, you can update the artifact lifetime.

Define artifact resolver locations (SAML)

When the artifact binding is enabled is enabled as one of the allowable bindings on the **Allowable SAML Bindings** screen, you must provide an artifact resolution service (ARS) endpoint. This is the location where PingFederate® sends back-channel requests to resolve artifacts received from the IdP.

SAML 2.0 connections allows multiple ARS endpoints. For SAML 1.x connection, you can only enter one ARS endpoint.

1. Enter an ARS endpoint.

- a) Enter the ARS endpoint URL.

You may enter a relative path (begin with a forward slash) if you have provided a base URL on the **General Info** screen.

If you are configuring a SAML 1.x connection, you can only enter one ARS endpoint on the **Artifact Resolver Location** screen.

- b) Optional: Enter an integer to the **Index** field for this ARS endpoint.

(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.

(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 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.



If you are editing an existing connection, you may reconfigure any ARS endpoint (or the source ID value for SAML 1.x).

Configure OpenID Provider information

1. On the **OpenID Provider Info** screen, provide the scopes, the endpoints, and the authentication scheme; for example:

 **Note:** If you have chosen to load the metadata from the OP on the **General Info** screen, the **Scopes** field and all endpoints are pre-populated (provided that the metadata contains the information).

Field	Description
Scopes	<p>The scopes to be included in the authentication and token requests to the OP. Multiple space-separated values are allowed.</p> <p>The default value (without loading metadata from the OP) is <code>openid</code>.</p> <p> 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.</p>
Authorization Endpoint	<p>The authorization endpoint at the OP.</p> <p>You may enter a relative path (begin with a forward slash) if you have provided a base URL on the General Info screen.</p> <p>There is no default value (without loading metadata from the OP).</p>
OpenID Connect Login Type	<p>The OpenID Connect client profile of the client. This client represents PingFederate and is created and managed at the OP.</p> <ul style="list-style-type: none"> If the client is configured to support the Basic Client profile, select Code. The resulting value of the <code>response_type</code> parameter is <code>code</code>. If the client is configured to support the Implicit Client profile, select Form POST. The resulting value of the <code>response_type</code> parameter is <code>id_token</code>. If the client is configured to support the Implicit Client profile <i>and</i> the target application requires the associated access token, select Form POST with access token. The resulting values of the <code>response_type</code> parameter are <code>id_token token</code>. <p>The default selection (without loading metadata from the OP) is Code.</p>
Authentication Scheme	<p>The client authentication method that PingFederate uses. Applicable and visible only to clients supporting the Basic Client profile.</p> <ul style="list-style-type: none"> Select Basic to submit credentials via HTTP Basic authentication. Select POST to submit credentials via POST. Select Private Key JWT to authenticate via the <code>private_key_jwt</code> Client Authentication method, see Client Authentication in the OpenID Connect specification (openid.net/specs/openid-connect-core-1_0.html#ClientAuthentication). <p>The default selection (without loading metadata from the OP) is Basic.</p>
Authentication Signing Algorithm	<p>Select the algorithm that PingFederate uses to sign the JWT.</p>

Field	Description
Token Endpoint, User Info Endpoint, and JWKS URL	<p>Applicable and visible only when Private Key JWT is the chosen authentication scheme.</p> <p> Note: Based on the chosen signing algorithm, PingFederate selects the signing JSON Web Key (JWK) from its JWK Set (JWKS) at runtime.</p> <p>In order for the OP to validate the signed JWT, 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 <code><Base URL>/pf/JWKS</code>, where Base URL is defined on the Server Configuration > Server Settings > Federation Info screen. For example, if the Base URL field value is <code>https://www.example.com</code>, the PingFederate JWKS URL is <code>https://www.example.com/pf/JWKS</code>. You can pass the JWKS URL directly to the OP or have the OP contact the PingFederate OpenID Connect Metadata endpoint to obtain the information (see OpenID Provider configuration endpoint on page 476).</p> <p>Various OAuth 2.0 and OpenID Connect 1.0 endpoints at the OP. For more information, see openid.net/connect.</p> <p>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.)</p> <p>The User Info Endpoint field is optional. If omitted, PingFederate only has access to the end-user claims from the ID tokens.</p> <p>The JWKS URL is required in order for PingFederate to validate the inbound ID tokens from the OP.</p> <p>There are no default values (without loading metadata from the OP).</p>
Sign Request	<p>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.</p> <p>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 http://openid.net/specs/openid-connect-core-1_0.html#RequestObject.</p> <p>This check box is not selected by default, in which case PingFederate sends request parameters via multiple query parameters, unsigned.</p>
Request Signing Algorithm	<p>Select the algorithm that PingFederate uses to sign the request object.</p> <p>Applicable and visible only when the Sign Request check box is selected.</p> <p> Note: PingFederate automatically selects the signing JSON Web Key (JWK) based on the selected signing algorithm from its JWK Set (JWKS).</p> <p>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 <code><Base URL>/pf/JWKS</code>, where Base URL is defined on the Server Configuration > Server Settings > Federation Info screen. For example, if the Base URL field value is <code>https://www.example.com</code>, the PingFederate JWKS URL is <code>https://www.example.com/pf/JWKS</code>. You can pass the JWKS URL directly to the OP or have the OP contact</p>

Field	Description
	the PingFederate OpenID Connect Metadata endpoint for it (see OpenID Provider configuration endpoint on page 476).

- Optional: Remain on the **OpenID Provider Info** screen and specify the request parameters that are allowed to be included in the authentication requests to the OP under **Request Parameters** (see [Configure request parameters and SSO URLs](#) on page 395).

Configure default target URLs

Use the **Protocol Settings > Overrides** screen to assign a default target URL for this IdP connection.

Optional: Enter a URL in the **Default Target URL** field.

If specified, the value overrides the default target setting defined on the **SP Configuration > Default URLs** screen.



Note: The SAML 1.x specifications for IdP-initiated 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 on this screen will not be used for those transactions.

Override authentication context in an IdP connection

Authentication context values can be mapped between the local and remote values in an OpenID Connect or a SAML 2.0 IdP connection. 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 may 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.

- Select the applicable IdP connection from the **SP Configuration** menu.

If the applicable connection is not one of the most recently edited connections, click **Manage All** under **IdP Connections**, and then select it from the **IdP Connections** screen.

You may also configure authentication context overrides when you create a new IdP connection.

- On the **Activation & Summary** screen, scroll down to the **Protocol Settings** section, and then click **Overrides**.
- On the **Overrides** screen, specify the **Local** and **Remote** entry, and then click **Add**.
- Repeat the previous step to define additional mappings.

Use the **Edit**, **Update**, and **Cancel** workflow to make or undo a change to an existing value. Use the **Delete** and **Undelete** workflow to remove an existing value or cancel the removal request.

- Click **Save** to complete the configuration.

Alternatively, click **Next** to carry on with the rest of the connection settings.

Example

Suppose you are the 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 IdP 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

Authentication method	AuthnContext values
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:KerberosAuthn	KerberosAuthn
urn:oasis:names:tc:SAML:2.0:ac:classes:unspecified	
urn:oasis:names:tc:SAML:2.0:ac:classes:unspecified	

The first entry maps KerberosAuthn 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.

Configure signature policy

The **Signature Policy** screen provides options controlling how digital signatures are used for SAML and WS-Federation SSO messages. The choices made on this screen depend on your partner agreement (see [Digital signing policy coordination](#) on page 77).

Digital signing is required for SAML response messages sent from the IdP via 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 screen. Note that 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 SP may 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** screen.)

- To continue, select the option (or options) based on your partner agreement.

If you are editing an existing connection, you can reconfigure the digital signature policy, which may require additional configuration changes in subsequent tasks.

Specify 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 may also agree to encrypt all or part of an assertion to improve privacy. If so, you can configure these settings on the **Encryption Policy** screen.

If the name identifier (or SAML_SUBJECT) of an assertion is encrypted, you and your partner may also want to encrypt the identifier in subsequent SAML 2.0 SLO messages (if you are using an SLO profile).

Note that “The entire assertion” selection on the Encryption Policy screen includes the SAML_SUBJECT and all attributes.

To enable encryption:

1. Select the **Allow encrypted SAML Assertions and SLO messages** option.
2. Choose whether this 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 SLO messages, allow its encryption for inbound messages, or both.

If you are editing an existing connection, you can reconfigure the XML encryption policy, which may require additional configuration changes in subsequent tasks.

Review protocol settings

When you are finished with the **Protocol Settings** task, you can review the configuration on its **Summary** screen.

- If you need to make any changes, click the heading over the information you want to edit.
- If you are creating a new connection and want to keep your configuration, click **Done**.
- If you are editing an existing configuration and want to keep your changes, click **Save**.

Review Browser SSO settings

When you are finished with the **Browser SSO** primary task, you can review the configuration on its **Summary** screen.

- If you need to make any changes, click the heading over the information you want to edit.
- If you are creating a new connection and want to keep your configuration, click **Done**.
- If you are editing an existing configuration and want to keep your changes, click **Save**.

Manage Attribute Query profile

At the Attribute Query step you configure your connection to request user attributes from your partner IdP, if you have chosen this option (see [Choose IdP connection options](#) on page 342). Attribute queries are not dependent on single sign-on but may be used independently or in conjunction with Browser SSO or provisioning to provide flexibility in how a user authenticates with SP applications (see [Attribute Query and XASP](#) in the “Supported Standards” chapter of the *Get started with PingFederate* guide).

Set the Attribute Authority Service URL

Attribute Authority is the term used to refer to an IdP that provides user attributes to an *Attribute Requester* (your SP site). The Attribute Authority Service URL corresponds to the endpoint location where Attribute Query requests are received by your IdP partner (see [Attribute Query and XASP](#) in the Supported Standards chapter of the *Get started with PingFederate* guide).

To configure the URL:

- Enter the fully qualified URL or a relative path if you have defined a base URL (see [Identify the partner](#) on page 344).

Map 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, then you need to map them on this screen. When the 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.

This information must be predetermined in your agreement with this connection partner.

To map attributes:

1. Enter the Local Name and Remote Name of an attribute and click **Add**.

Repeat this step for all attributes requiring mapping.

2. Click **Next**.

To edit a mapping:

1. Click **Edit** under Action for the mapping.
2. Make your change(s) and click **Update**.



Note: If you change your mind, ensure that you click the **Cancel link** in the Actions column, not the **Cancel button**, which discards any other changes you might have made in this configuration.

3. Click **Done** and then **Save** on the Attribute Query screen.

Configure security policy for Attribute Query

This screen allows you to specify the digital signing and encryption policy to which you and your partner have agreed. These selections will trigger requirements for setting up Credentials (see [Configure security credentials](#) on page 385).

This screen also allows you to mask incoming attribute values in log files (see [Attribute masking](#) on page 74). When you enable this selection, all user attributes returned from this IdP are masked.

To configure attribute-query security policy for this partner:

- Check or clear the check boxes and click **Next** or **Done**.

Review the Attribute Query settings

On the Summary screen you can review the Attribute Query configuration.

To reconfigure saved profiles:

1. Click the heading over the information you want to change.
2. Click **Done** on the screen containing your change.

If you need to make additional changes, do so and continue by clicking **Done** until you reach the Attribute Query screen.

3. Click **Save** on the Attribute Query screen.

Configure just-in-time provisioning

PingFederate[®]'s Just-in-Time (JIT) Provisioning allows SPs to create user accounts on the fly during SSO events, based on attributes received in SSO tokens from IdPs. An SP can also use the feature to update existing user records.



Note: This configuration task is presented in the administrative console only when the **JIT Provisioning** check box is selected in the **Connection Options** screen.

1. Create a new IdP connection or select an existing IdP connection on the **SP Configuration** menu.
2. On the **Connection Type** screen, select the **Browser SSO Profiles** check box and a protocol from the list.
3. On the **Connection Options** screen, select the **Browser SSO** check box and then the **JIT Provisioning** check box.
4. Complete the **Browser SSO** configuration.
5. On the the **JIT Provisioning** screen, click **Configure User Provisioning** to begin the configuration of JIT provisioning.

Select 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 IdP in a follow-on Attribute Query transaction. The **User Attributes** screen displays the attributes expected in the assertion from this IdP.



Note: Attribute Query is a SAML 2.0 profile. For OpenID Connect, SAML 1.x, and WS-Federation connections, this screen is not presented; PingFederate® uses only attributes from the assertion for user provisioning.

- If you and your IdP partner have agreed to use the Attribute Query profile for provisioning, select that option before leaving this screen.

You configure the Attribute Query profile later in the task flow.

Identify the user repository

PingFederate®'s JIT Provisioning currently supports LDAP v3-compliant user stores and the JDBC-compliant Microsoft SQL Server.



Note: We recommend using the latest version of the Microsoft SQL Server JDBC Driver (version 4.1 or higher), which is compatible with current and recent Server JRE versions.

- On the **User Repository** screen, select a data store from the **Active Data Store** list.



Tip: If the desired data store is not shown in the list, then PingFederate has not yet been configured to access it. Click **Manage Data Stores** to configure your data store.

- If you are using an LDAP store, refer to the sections immediately following:
 - [Specify an LDAP user-record location](#) on page 367
 - [Enter an LDAP filter](#) on page 368
 - [Identify provisioning attributes for LDAP](#) on page 368
- If you are using a Microsoft SQL Server, skip to this section:
 - [Choose a SQL method](#)

Specify an LDAP user-record location

After choosing a data store, indicate where in the store PingFederate® should write new user records or update existing ones. Start by specifying where user records are located in your data store.



Note: This screen appears only when an LDAP data store is chosen on the **User Repository** screen.

- Enter the base DN in the text field.

A base distinguished name (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.

Enter an LDAP filter

On the **Unique User ID** screen, create an LDAP filter to identify user accounts to be provisioned (or updated) during SSO events. PingFederate® uses this expression in conjunction with the Base DN value defined on the **Location** screen to locate existing account records and to add new ones.

Note: This screen appears only when an LDAP data store is chosen on the **User Repository** screen.

- Enter the statement in the **Filter** text field.

The filter is in the form: `attribute=${value}`

Note: 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-data store. Click the link near the lower-left corner of the screen to see a list of available attributes.

The right side of the filter generally 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**.

Tip: If you are unfamiliar with writing LDAP queries, please refer to the documentation accompanying your LDAP installation.

Identify provisioning attributes for LDAP

On the **Attributes** screen, select the data-store attributes to be provisioned.

Note: This screen appears only when an LDAP data store is chosen on the **User Repository** screen.

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.

Choose a SQL method

For JDBC data stores, PingFederate® allows you to map attributes directly to a single database table (the default) or to SQL stored-procedure parameters.

Note: This screen appears only when a Microsoft SQL Server data store is chosen on the **User Repository** screen.

- Make a selection as needed and click **Next**.

Depending on the selection, different steps appear under the **JIT Provisioning** task. Refer to the sections indicated below for more information:

- If you are mapping attributes directly to a table, refer to the topics sections immediately following:
 - [Specify a database user-record location](#) on page 369
 - [Specify a unique-ID database column](#) on page 369
- If you are using a stored procedure, skip to this section:
 - [Specify a stored-procedure location](#) on page 369

Specify a database user-record location

For database provisioning to a table, indicate where PingFederate® should write new user records or update existing ones.



Note: This screen appears only when a Microsoft SQL Server data store is chosen on the **User Repository** screen and the **Table** option is selected on the **SQL Method** screen.

- 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.
Columns to fulfill	For the selected table, all attributes and their data types are displayed. If the list of columns is not or might not be current (due to any recent changes), click Refresh . Later on the Attribute Fulfillment screen, all attributes must be mapped for the database insertion to succeed, although a null entry may be used for optional attributes.



Tip: Click the link near the lower-left corner of the screen to see a list of available attributes from the SSO token.

Specify a unique-ID database column

PingFederate uses the column specified on this screen to check whether a user record already exists for the incoming SSO token.





Note: This screen appears only when a Microsoft SQL Server data store is chosen on the **User Repository** screen and the **Table** option is selected on the **SQL Method** screen.

- Select a column that represents a unique characteristic about the database entry for a particular user.

Specify a stored-procedure location


If you are using a stored procedure for provisioning the user database, specify its location on this screen.

 **Note:** This screen appears only when a Microsoft SQL Server data store is chosen on the **User Repository** screen and the **Stored Procedure** option is selected on the **SQL Method** screen.

 **Important:** The database account used by PingFederate must have access to the schema in which the stored procedure is located and to execute permission for the procedure.

- Select the database schema and the stored procedure.

Field	Description
Schema	Select the table structure that contains stored procedure within the database.
Stored Procedure	Select the stored procedure needed to provision the user database.
Procedure parameters to fulfill	For the selected procedure, all parameters and their data types are displayed. If the list of parameters is not or might not be current (due to any recent procedure revisions), click Refresh . You map attribute values from the SSO token to the parameters on the Attribute Fulfillment screen.

 **Tip:** Click the link near the lower-left corner of the screen to see a list of available attributes from the SSO token.

Map attributes to a user account

The **Attribute Fulfillment** screen provides a means of mapping the 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. Besides values obtained from the SSO token, you may also map from the context of the SSO token, text with or without references values from the SSO token, and expression (if enabled).

If a Microsoft SQL Server data store is chosen on the **User Repository** screen, the **Attribute Fulfillment** screen, you may also 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 may use a PingFederate® Java conversion method via OGNL expressions.


1. Select a source from the list for each target attribute or parameter.

- Assertion or Provider Claims

Values are contained in the SSO token from this IdP. When you make this selection, the associated **Value** list is populated by the attribute contract).

- Context

Values are returned from the context of the transaction at runtime.

 **Note:** The HTTP Request is retrieved as a Java objects 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.)

- Attribute Query


This choice appears only if you have chosen to use the Attribute Query profile for provisioning.

To map an attribute-query value, use this syntax: `${query_attribute}`

You can also 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:** If you have not already done so, you may enable OGNL expression by editing the `org.sourceid.common.ExpressionManager.xml` file in the `<pf_install>/pingfederate/server/default/data/config-store` directory. 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 on the **Server Configuration > Cluster Management** screen, and restart all nodes.

This option provides more complex mapping capabilities; for example, transforming incoming values into different formats. All of the variables available for text entries (see below) are also available for expressions.

 **Tip:** If multiple attribute values from one or multiple sources need to be mapped 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** for such attributes, click **Datetime OGNL Examples** under the text box for syntax information and examples.

- System Managed (if applicable)

This mapping option appears only when any automatically assigned attributes are among columns to be provisioned; for example, 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 text box 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 text box to explicitly set `null` in the column.

2. Select (or enter) an attribute value.

All values must be mapped. However, for optional table columns, you may leave a text box blank (or, for string data types, enter `null` to avoid empty strings).

Note that no value is required for **System Managed** attributes.

 **Note:** For Active Directory, enter `user` in the text box for objectClass. For the Oracle Directory Server, enter `inetOrgPerson`.

3. Optional: When mapping to a Microsoft SQL Server data store, test the insertion.

SQL method

Steps

Table

1. Click **Test insert into 'table'**.
2. Enter values for each applicable target parameter.
3. Click **Test Insert**.

If the test succeeds, a confirmation is displayed along with the values inserted.



Caution: Unless you wish to keep the test values in the database, click **Roll Back All Test Inserts**.

Stored procedure

1. Click **Test call to 'procedure'**.
2. Enter values for each applicable target parameter.

SQL method

Steps

3. Click **Test Stored Procedure Call**.

For stored procedures, only a confirmation is displayed 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, if needed, must be handled manually.

When finished, click **Return to Attribute Fulfillment**.

Choose an event trigger

On the **Event Trigger** screen, choose whether PingFederate® initiates user provisioning only when the user identifier is new or every time your site receives an SSO token. In the latter case (for all SSO tokens), an existing user account is always updated with incoming attributes.



Note: This screen does not appear for a Microsoft SQL Server data store 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 (if desired).

Configure an error handling method

If user provisioning fails for any reason during SSO events, you can choose to stop the SSO or continue the process by passing the attributes on to your target application.

When SSO is aborted, the user is redirected to an error page and the failure is written to the server log.

Review the JIT provisioning configuration

The **Summary** screen provides an overview of your provisioning configuration.

1. Click the heading over the information you want to change.
2. Click **Done** on the screen containing your change.

If you need to make additional changes, do so and continue by clicking **Done** until you reach the **JIT Provisioning** screen.

3. Click **Save** (or **Next**) on the **Inbound Provisioning** screen.

Configure SCIM inbound provisioning

This configuration provides for two-way mapping of attributes. The first facilitates SCIM operations used to create and update records in the data store (see [Write user information to the data store](#) on page 377). The second allows the same SCIM client to retrieve those records and have the attribute values mapped back to their corresponding designation in the client store (see [Configure a SCIM response](#) on page 379).

The dual mapping is intended to provide greater flexibility, especially when needed for OGNL-expression transformations; for example, converting two attributes into one multi-valued attribute and then back again.



Note: SCIM-client requests must include authentication credentials, which you configure later on the **Credentials > Back-Channel Authentication** screen. The same credentials needed for SSO or other types of transactions enabled as part of this IdP connection, if configured, are also used for SCIM transactions.

1. If you have not already done so, enable the **Inbound Provisioning** option (under the Server Provider role) on the **Server Configuration > Server Settings > Roles & Protocols** screen.
2. Create a new IdP connection or select an existing IdP connection on the **SP Configuration** menu.
3. On the **Connection Type** screen, select the **Inbound Provisioning** check box and one of these two options: **User Support** or **User and Group Support**.
4. On the **Inbound Provisioning** screen, click **Configure Inbound Provisioning** to begin the configuration of SCIM inbound provisioning.

Specify the user repository

PingFederate[®] currently supports AD user stores (requires an LDAPS connection to the data store) and custom identity store provisioners for Inbound Provisioning.

- Choose either the **Active Directory Data Store** or **Identity Store Provisioner** option and then specify the data store from the list.



Tip: If the correct data store is not shown in the list, then PingFederate is not yet configured to access the store. Click **Manage ...** to set up the desired data store.

Identify an LDAP user-record location

After choosing a data store, indicate where in the data store user and group records exist so PingFederate[®] can create, read, update, or delete/disable them.




Note: This screen appears only if you are configuring an LDAP user store for provisioning.

- Enter the base distinguished name of the tree structure where user records are stored in the **Base DN** field. PingFederate looks only at this node level or below it for user accounts that need provisioning.

Define a unique ID

On the **Unique ID** screen, create an LDAP filter to resolve user accounts for SCIM operations. PingFederate[®] uses this expression in conjunction with the **Base DN** value (defined on the **Location** screen) to add new account records.



Note: This screen appears only if you are configuring an LDAP user store for provisioning.

- Enter the statement in the **Filter** text field.

The filter is in the form: `attribute=${value}`

Note: Unlike filters used to retrieve LDAP attributes for adapter mapping, do not enclose the statement in parentheses.

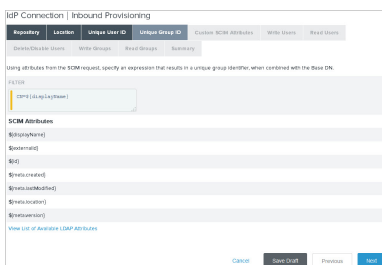
The left-side variable is an attribute in your user-data store. Click the link near the lower-left corner of the screen to see a list of available attributes.

The right side of the filter generally uses one or more attribute values passed in from the SCIM request. Variables for these attributes, including the correct syntax, are listed under **SCIM Attributes**.

Tip: If you are unfamiliar with writing LDAP queries, please refer to the documentation accompanying your LDAP installation.

Define a unique group ID

On the **Unique Group ID** screen, create an LDAP filter to resolve groups for SCIM operations. PingFederate® uses this expression in conjunction with the **Base DN** value (defined on the **Location** screen) to add new groups.



Note: This screen appears only if you are configuring an LDAP user store for provisioning and the **User and Group Support** option is selected on the **Connection Type** screen.

- Enter the statement in the **Filter** text field.

The filter is in the form: `attribute=${value}`

Note: 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-data store. Click the link near the lower-left corner of the screen to see a list of available attributes.

The right side of the filter generally uses one or more attribute values passed in from the SCIM request. Variables for these attributes, including the correct syntax, are listed under **SCIM Attributes**.

Tip: If you are unfamiliar with writing LDAP queries, please refer to the documentation accompanying your LDAP installation.

Define custom SCIM attributes

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** screen.

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 Multi-Valued Attributes
- Complex Attributes
- Complex Multi-Valued 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 Multi-Valued	Not Applicable
options	Complex	quantity, strike, first, and last
tablets	Complex Multi-Valued	model, serial, and type.



Note: type is a reserved sub-attribute for a complex multi-valued attribute.



Tip: For more information about SCIM and attribute types, see the website www.simplecloud.info.

- To specify a schema extension, enter the URI of the schema extension in the **Extension Namespace** field.



Tip: 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.

The administrative console displays the **Custom SCIM Attribute Options** screen, where you may

- Change the attribute name
- Set the attribute as a simple multi-valued attribute
- Add sub-attributes to make it a complex attribute
- Add sub-attributes and set the attribute as a complex multi-valued attribute

(For more information, see [Configure custom SCIM attribute options](#) on page 376.)

Configure custom SCIM attribute options

For the chosen custom attribute, use the **Custom SCIM Attribute Options** screen to change the attribute name, set the attribute as a simple multi-valued attribute, add sub-attributes to make it a complex attribute, and add sub-attributes and set the attribute as a complex multi-valued attribute.

- 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 multi-valued 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 click **Add**.



Tip: Use the **Edit/Update/Cancel** links to make or undo a change to the name of a sub-attribute. Use the **Delete/Undelete** links to remove a sub-attribute or cancel the deletion.

Repeat this step to add more sub-attributes as needed and then click **Done**.

Custom SCIM Attribute Configuration. You are presented with options to make the attribute complex and/or multi-valued.

Sub-Attributes	Action
identity	Edit Delete
email	Edit Delete
phone	Edit Delete
url	Edit Delete
	Add

Is Multi-Valued

- To define the custom attribute as a complex multi-valued attribute:

1. Enter a sub-attribute and click **Add**.



Tip: Use the **Edit/Update/Cancel** links to make or undo a change to the name of a sub-attribute. Use the **Delete/Undelete** links to remove a sub-attribute or cancel the deletion.

Repeat this step to add more sub-attributes as needed.

2. Select the **Is Multi-Valued** check box.

3. If you have chosen Active Directory as your user store (see [Specify the user repository](#) on page 373), you must specify at least one value under the **Types** column for `type`, a reserved sub-attribute for a complex multi-valued attribute.

4. Click **Done**.

Custom SCIM Attribute Configuration. You are presented with options to make the attribute complex and/or multi-valued.

Sub-Attributes	Action
identity	Edit Delete
email	Edit Delete
phone	Edit Delete
url	Edit Delete
	Add

Is Multi-Valued

Types	Action
type	Edit Delete
	Add

Write user information to the data store

To configure how PingFederate® completes create and update operations for user accounts from a SCIM request, identify incoming attributes and map them to data-store attributes.

Write Users Configuration

Configure Write Users

- Click **Configure Write Users** to continue.

Identify inbound provisioning attributes for LDAP

On the **Attributes** screen, select the data-store attributes you want to provision.

Select attributes in the user repository that should be populated during a user provisioning request.

Root Object Class	Attribute	Action
	email	Remove
	phone	Remove
	url	Remove

Show All Attributes: Enabled



Note: This screen appears only if you are configuring an LDAP user store for provisioning.

Several 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** screen. In this case, the values you supply are used. The objectGUID and userAccountControl attributes cannot be overridden and are ignored if selected.

1. Select a root object class and an attribute from the lists, then click **Add Attribute**.



Important: Do *not* add cn as one of the attributes.

2. Repeat the previous step for each attribute requiring provisioning.

Map attributes to user accounts

Use this screen to map attribute values in the SCIM request to user-account attributes.

1. Select a source from the list for each target attribute.

- **Context**

When selected, the **Value** list is populated with the available context of the transaction. Select the desired context from the list.



Note: The **HTTP Request** context value is retrieved as a Java object rather than text. For this reason, OGNL expressions are more appropriate to evaluate and return values (see **Expression**).

- **Expression**



Tip: If you have not already done so, you may enable OGNL expression by editing the `org.sourceid.common.ExpressionManager.xml` file in the `<pf_install>/pingfederate/server/default/data/config-store` directory. 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 on the **Server Configuration > Cluster Management** screen, and restart all nodes.

This option provides more complex mapping capabilities; for example, transforming incoming values into different formats. All of the variables available for text entries (see below) 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 drop-down list is populated by defined components of the SCIM request.

- **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. Select (or enter) an attribute value.

All target attributes must be mapped.

3. Click **Done**.


Review user mapping (Write Users) configuration

The **Summary** screen provides an overview of the inbound provisioning configuration for request mapping.

- When you finish with a new configuration, click **Done**.

Configure a SCIM response


To configure a SCIM response to a request to read and return provisioned SCIM attributes, identify and map the user account attributes you want to include.



- Click **Configure Read Users** to continue.

Identify expected user attributes for the SCIM response

An attribute contract is a set of user attributes that you and your partner have agreed will be sent in a SCIM response for this connection. The attributes you mapped to user account attributes in the **Write Users** flow appear under **Attribute Contract**.



Click **Available SCIM Attributes** near the lower-left corner of the screen 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
- active
- 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**.

- To modify an attribute name or masking selection, click **Edit** under **Action** for the attribute, make the change, and click **Update**.




Note: If you change your mind, ensure that you click **Cancel** under **Actions**, not the **Cancel** button, which discards any other changes you might have made in the configuration steps.

- To delete an attribute, click **Delete** under **Action** for the attribute.

Identify LDAP attributes for the SCIM response

Select the LDAP attributes you want to map to attributes in the SCIM response.

 **Note:** This screen appears only if you are configuring an LDAP user store for provisioning.

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.

Map attributes into the SCIM response

Use this screen to map outgoing user-account attributes to SCIM responses to READ requests.

1. Select a source from the list for each target attribute.

- **Context**

When selected, the **Value** list is populated with the available context of the transaction. Select the desired context from the list.



Note: The **HTTP Request** context value is retrieved as a Java object rather than text. For this reason, OGNL expressions are more appropriate to evaluate and return values (see **Expression**).

- **Expression**



Tip: If you have not already done so, you may enable OGNL expression by editing the `org.sourceid.common.ExpressionManager.xml` file in the `<pf_install>/pingfederate/server/default/data/config-store` directory. 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 on the **Server Configuration > Cluster Management** screen, and restart all nodes.

This option provides more complex mapping capabilities; for example, transforming outgoing values into different formats. All of the variables available for text entries (see below) are also available for expressions.



Tip: 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 is populated by the LDAP attributes you identified for this data store.

- **Identity Store**

Values are returned from your query. When you make this selection, the Value list is populated by the Identity Store attributes you identified for this data store.

- **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. Select (or enter) an attribute value.

All target attributes must be mapped.

3. Click **Done**.

Review SCIM response (Read Users) configuration

The **Summary** screen provides an overview of the inbound provisioning configuration for SCIM response mapping.

- When you finish with a new configuration, click **Done**.

Configure the handling of SCIM delete requests

Use this screen to define how SCIM delete requests are handled within your user data store.

- **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 data store.

- **Note:** This screen appears only if you are configuring an LDAP user store for provisioning.

- Select **Disable User** to make the user inactive within the data store. This approach might be 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 (see the [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.
- Select **Permanently Delete User** to remove the user from the data store.

Write group information to the data store

To configure how PingFederate® completes create and update operations for groups from a SCIM request, identify incoming attributes and map them to data-store attributes.

- Click **Configure Write Groups** to continue.

Identify inbound provisioning group attributes for LDAP

On the **Attributes** screen, select the data-store attributes you want to provision.

- **Note:** This screen appears only if you are configuring an LDAP user store for provisioning and the **User and Group Support** option is selected on the **Connection Type** screen.

Several attributes are managed internally by PingFederate and do not require mapping:

- objectClass

- objectGUID
- member

You can override the internal management of objectClass by selecting and mapping it to a SCIM attribute on the **Attribute Fulfillment** screen. In this case, the values you supply are used. The objectGUID and member attributes cannot be overridden and are ignored if selected.

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.

Map attributes to groups

Use this screen to map attribute values in the SCIM request to group attributes.

1. Select a source from the list for each target attribute.

- **Context**

When selected, the **Value** list is populated with the available context of the transaction. Select the desired context from the list.



Note: The **HTTP Request** context value is retrieved as a Java object rather than text. For this reason, OGNL expressions are more appropriate to evaluate and return values (see **Expression**).

- **Expression**



Tip: If you have not already done so, you may enable OGNL expression by editing the `org.sourceid.common.ExpressionManager.xml` file in the `<pf_install>/pingfederate/server/default/data/config-store` directory. 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 on the **Server Configuration > Cluster Management** screen, and restart all nodes.

This option provides more complex mapping capabilities; for example, transforming incoming values into different formats. All of the variables available for text entries (see below) 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 Group**

When you make this selection, the associated Value drop-down list is populated by defined components of the SCIM request.

- **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. Select (or enter) an attribute value.

All target attributes must be mapped.

3. Click **Done**.


Review group mapping (Write Groups) configuration

The **Summary** screen provides an overview of the inbound provisioning configuration for request mapping.

- When you finish with a new configuration, click **Done**.

Configure a SCIM response for groups


To configure a SCIM response to a request to read and return provisioned SCIM attributes, identify and map the group attributes you want to include.



- Click **Configure Read Groups** to continue.

Identify expected group attributes for the SCIM response

An attribute contract is a set of group attributes that you and your partner have agreed will be sent in a SCIM response for this connection. The attributes you mapped to group attributes in the Write Groups flow appear at the top of the screen.



Click **Available SCIM Attributes** near the lower-left corner of the screen 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
- 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**.

- To modify an attribute name or masking selection, click **Edit** under **Action** for the attribute, make the change, and click **Update**.




Note: If you change your mind, ensure that you click **Cancel** under **Actions**, not the **Cancel** button, which discards any other changes you might have made in the configuration steps.

- To delete an attribute, click **Delete** under **Action** for the attribute.

Identify LDAP group attributes for the SCIM response

Select the LDAP attributes you want to map to attributes in the SCIM response.




Note: This screen appears only if you are configuring an LDAP user store for provisioning and the **User and Group Support** option is selected on the **Connection Type** screen.

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.

Map group attributes into SCIM response

Use this screen to map outgoing group attributes to SCIM responses to READ requests.



1. Select a source from the list for each target attribute.

- **Context**

When selected, the **Value** list is populated with the available context of the transaction. Select the desired context from the list.



Note: The **HTTP Request** context value is retrieved as a Java object rather than text. For this reason, OGNL expressions are more appropriate to evaluate and return values (see **Expression**).

- **Expression**



Tip: If you have not already done so, you may enable OGNL expression by editing the `org.sourceid.common.ExpressionManager.xml` file in the `<pf_install>/pingfederate/server/default/data/config-store` directory. 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 on the **Server Configuration > Cluster Management** screen, and restart all nodes.

This option provides more complex mapping capabilities; for example, transforming outgoing values into different formats. All of the variables available for text entries (see below) are also available for expressions.



Tip: 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 is populated by the LDAP attributes you identified for this data store.

- **Identity Store**

Values are returned from your query. When you make this selection, the Value list is populated by the Identity Store attributes you identified for this data store.

- **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. Select (or enter) an attribute value.

All target attributes must be mapped.

3. Click **Done**.

Review SCIM response for groups (Read Groups) configuration

The **Summary** screen provides an overview of the inbound provisioning configuration for SCIM response mapping.

- When you finish with a new configuration, click **Done**.

Review the inbound provisioning configuration

On the **Summary** screen you can review the **Inbound Provisioning** configuration.

1. Click the heading over the information you want to change.

2. Click **Done** on the screen containing your change.

If you need to make additional changes, do so and continue by clicking **Done** until you reach the **Inbound Provisioning** screen.

3. Click **Save** (or **Next**) on the **Inbound Provisioning** screen.

Configure security credentials

The **Credentials** screen 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. Your connection configuration may involve any or all of the following:

- [Manage back-channel authentication](#) on page 385
- [Manage digital signature settings](#) on page 387
- [Manage signature verification settings](#) on page 387
- [Choose an encryption certificate \(SAML 2.0\)](#) on page 388
- [Choose a decryption key \(SAML 2.0\)](#) on page 389
- To continue, click **Configure Credentials**.

Manage back-channel authentication

When you configure 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 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.

Refer to the following topics for configuration steps:

- [Configure back-channel authentication for outbound messages](#) on page 385
- [Configure back-channel authentication for inbound messages](#) on page 386

Configure back-channel authentication for outbound messages

1. On the **Back-Channel Authentication** screen, click **Configure** to the right of the list of messages under **Send to your partner**.
2. On the **Outbound SOAP Authentication Type** screen, 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)** screen.

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** screen. If you have not yet created or imported the client certificate, click **Manage Certificates** to do so (see [Manage SSL client keys and certificates](#) on page 182).

 **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 a subsequent screen, **Digital Signature Settings**.

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 may be used in any combination or independently.

3. On the **Summary** screen for outbound authentication requirements:


- If you need to make any changes, click the heading over the information you want to edit.
- If you are creating a new connection and want to keep your configuration, click **Done**.
- If you are editing an existing configuration and want to keep your changes, click **Save**.

Configure back-channel authentication for inbound messages

1. On the **Back-Channel Authentication** screen, click **Configure** to the right of the list of messages under **Received from your partner**.
2. On the **Inbound Authentication Type** screen, 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)** screen.

 **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 screens (see next step).

Digital Signature (Browser SSO profile only)

You select a signing certificate on a subsequent screen, **Signature Verification Settings**.

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 may 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 may 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** screen.

Anchored

The partner certificate must be signed by a trusted certificate authority (CA). Optionally, you may 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 **Server Configuration > Trusted CAs** screen.


Unanchored

The partner certificate is self-signed or you want to trust a specified certificate.



Note: When anchored certificates are used between partners, certificates may 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 77.)

Trust model	Subsequent steps
Anchored	<p>On the Subject DN screen:</p> <ol style="list-style-type: none"> 1. Enter the Subject DN of the certificate. 2. (Optional) Select the Restrict Issuer check box and enter the Issuer DN of the certificate. <p> 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.</p>
Unanchored	<p>On the SSL Verification Certificate screen, select the client certification from your partner.</p> <p>If you have not yet imported the client certificate from your partner, click Manage Certificates to do so (see Manage certificates from partners on page 190).</p>

4. On the **Summary** screen for inbound authentication requirements:

- If you need to make any changes, click the heading over the information you want to edit.
- If you are creating a new connection and want to keep your configuration, click **Done**.
- If you are editing an existing configuration and want to keep your changes, click **Save**.

Manage digital signature settings

This step defines the private key you will use to sign SSO authentication or attribute requests (optionally) or SAML 2.0 SLO messages for this IdP. In addition, the step allows you to include "Key Info" with the XML message if you and your partner have agreed to this option.

Digital signing applies to SP-initiated SSO under SAML 2.0, when specified by your partner agreement, and to either SLO profile using the POST or redirect bindings. The step 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.)

1. Select a signing certificate from the list.

If you have not yet created or imported your certificate into PingFederate, click **Manage Certificates** (see [Manage digital signing certificates and decryption keys](#) on page 184).

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.

Manage signature verification settings


Under SAML 2.0 specifications, when your site receives any SAML 2.0 messages via the POST or Redirect bindings, the messages must be digitally signed. 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 IdP, SSO assertions, SAML 2.0 artifacts, or SOAP messages might also require signatures.

1. On the **Signature Verification Settings** screen, click **Manage Signature Verification Settings**.
2. On the **Trust Model** screen, select a trust model on the **Certificate Verification Method** screen.


Anchored

The partner certificate must be signed by a trusted certificate authority (CA). Optionally, you may 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 **Server Configuration > Trusted CAs** screen.



 **Important:** If you are using the redirect binding for SLO, you cannot use anchored certificates because SAML 2.0 does not permit certificates to be included using this transport method.

Unanchored

The partner certificate is self-signed or you want to trust a specified certificate.

 **Note:** When anchored certificates are used between partners, certificates may 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 77.)

Trust model	Subsequent steps
Anchored	<p>On the Subject DN screen:</p> <ol style="list-style-type: none"> 1. Enter the Subject DN of the certificate or extract it from your SP partner's certificate if the certificate is stored on an accessible file system. 2. (Optional) Select the Restrict Issuer check box and enter the Issuer DN of the certificate. Alternatively, extract it from your partner's certificate. <p> 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.</p>
Unanchored	<p>On the Signature Verification Certificate screen:</p> <ol style="list-style-type: none"> 1. Select a primary certificate from the list. <p>If you have not yet imported the certificate from your partner, click Manage Certificates to do so (see Manage certificates from partners on page 190).</p> <ol style="list-style-type: none"> 2. Select a secondary certificate from the list. <p> Tip: Use this field if your partner has sent you a new certificate to replace one that is ready to expire. The server will automatically verify against the secondary certificate when the primary one expires.</p>

3. On the **Summary** screen for signature verification:

- If you need to make any changes, click the heading over the information you want to edit.
- If you are creating a new connection and want to keep your configuration, click **Done**.
- If you are editing an existing configuration and want to keep your changes, click **Save**.

Choose 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 SLO requests from your site may also be encrypted (if the connection uses SP-initiated SLO). For more information, see [Specify XML encryption policy \(for SAML 2.0\)](#) on page 364.


You must also choose a certificate if encryption of the name identifier is required for an Attribute Request profile.

1. Optional: Change the default settings under **Block Encryption Algorithm**.

Due to import control restrictions, the standard JRE distribution supports strong but not unlimited encryption. To use the strongest AES encryption, when permissible, download and install the appropriate version of "Java

Cryptography Extension (JCE) Unlimited Strength Jurisdiction Policy Files" from the [Oracle Downloads](http://www.oracle.com/technetwork/java/javase/downloads/index.html) website (www.oracle.com/technetwork/java/javase/downloads/index.html).

For more information about XML block encryption and key transport algorithms, see [XML Encryption Syntax and Processing from W3C](http://www.w3.org/TR/xmlenc-core/) (www.w3.org/TR/xmlenc-core/)

 **Note:** Under **Key Transport Algorithm, RSA-v1.5** is disabled for new connections for security reasons. If you are updating an existing connection that uses RSA-v1.5, consider changing the selection for increased security.

Note that the JCE provider libraries distributed with SafeNet Luna SA HSM does not support the RSA-OAEP algorithm at the time of the PingFederate 7.1 R2 release. As a result, RSA-v1.5 is still allowed when deploying PingFederate with a Luna SA HSM.

2. Select a partner certificate from the list.

If you have not yet imported the certificate from your partner, click **Manage Certificates** to do so (see [Manage certificates from partners](#) on page 190).

Choose 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 (see [Specify XML encryption policy \(for SAML 2.0\)](#) on page 364).

1. Select the primary XML decryption key from the list.

If you have not yet created or imported your certificate into PingFederate, click **Manage Certificates** (see [Manage digital signing certificates and decryption keys](#) on page 184).

2. Optional: Select the secondary XML decryption key from the list.


Review credential settings

When you finish configuring various settings under the **Credentials** task, you can review the configuration on its **Summary** screen.

- If you need to make any changes, click the heading over the information you want to edit.
- If you are creating a new connection and want to keep your configuration, click **Done**.
- If you are editing an existing configuration and want to keep your changes, click **Save**.

Review 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** screen.

 **Important:** When creating a new connection, the default connection status is **Inactive** when you reach the **Activation & Summary** screen.

Regardless of whether you choose to activate a new connection now or later, you must click **Save** on the **Activation & Summary** screen if you want to keep the new connection.

The SSO application endpoint (underneath the **Connection Status** field) is a sample URL that the developers of the target application might use to invoke SSO for the connection (see [View SP application endpoints](#) on page 336).

If you have selected the **No Mapping** option on the **Identity Mapping** screen, the **Summary & Activation** screen does not show the **SSO Application Endpoint** sample URL.

- If you need to make any changes, click the heading over the information you want to edit.
- If you are creating a new connection and want to keep your configuration, you must click **Save**.
- If you are modifying an existing connection (including toggling the connection status) and want to keep your changes, you must click **Save**.

OpenID Connect Relying Party support

PingFederate® is capable of leveraging identities from OpenID Providers (OPs) to complete Browser 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. In essence, PingFederate retrieves identity information from the OP and passes the end-user claims, which are basically 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, a modern standard that has been gaining momentum in the industry.

If the OP supports the OpenID Connect Discovery specification, the connection setup can be 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.

In addition, the protocol settings can be refreshed 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 could be mapped to the target applications later. In the rare 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 may 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 may define additional request parameters, which could be included in the authentication requests to support OP-specific use cases. Administrators may restrict the values to those defined in the configuration. Alternatively, administrators may allow the target applications to optionally override the values at runtime. Furthermore, 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 via multiple query parameters, unsigned. Optionally, administrators can configure PingFederate to send request parameters via a request object by value, in which case request parameters are represented by individual claims in a signed 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 http://openid.net/specs/openid-connect-core-1_0.html#RequestObject.

Processing steps

1. A user starts a Browser SSO request at the `/sp/startSSO.ping` SP application endpoint or the `/sp/init_login.ping` SP protocol endpoint.
2. PingFederate (the RP) sends to the OP via the browser an authentication requests 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) via 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) via the browser.

The access token (if any) may also be passed to the target application, so that API security use cases can be layered on top of the browser-based SSO request.

Note that if the User Info 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 <https://openid.net/connect>.

Create an OpenID Connect IdP connection

1. If you have not already done so, enable the OpenID Connect protocol for IdP connections.
 - a) Go to the **Server Configuration > Server Settings > Roles & Protocols** screen.
 - b) Ensure the SP role is selected.
 - c) Select the **OpenID Connect** check box to activate the protocol for IdP connections.
2. On the **SP Configuration** menu, create a new IdP connection.
3. On the **Connection Type** screen, select the **Browser SSO Profiles** check box and select **OpenID Connect** from the **Protocol** list.

Note that when OpenID Connect is the chosen protocol, the other types become unavailable.

4. On the **Connection Options** screen, you may enable JIT provisioning, OAuth attribute mapping (which requires the OAuth 2.0 Authorization Server role), or both.

For simplicity, this topic focuses only on managing OpenID Connect IdP connection settings.

5. On the **General Info** screen:

- 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, please refer to 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 (see [step 12](#)).

- b) Optional: Click **Load Metadata**.



Tip: Loading metadata from the OP expedites the connection setup. You may also update an existing connection by reloading metadata.

6. On the **Browser SSO** screen, click **Configure Browser SSO**.
7. On the **User-Session Creation** screen, click **Configure User-Session Creation**.
8. On the **Identity Mapping** screen, you have three choices:
 - 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.




Tip: End-user claims are basically user attributes found in ID tokens or obtained from the User Info endpoint at the OP.

For illustration, this topic uses the **Account Mapping** configuration.

9. On the **Attribute Contract** screen, 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** screen, the attribute contract is populated automatically.

10. On the **Target Session Mapping** screen, click **Map New Adapter Instance** to map end-user claims to the target application through an SP adapter instance.

(You may also map an authentication policy contract if your PingFederate server is a federation hub that bridges an OP to an SP.)

Follow the administrative console to fulfill the SP adapter contract (or the authentication policy contract). Like other IdP connections, you have the options to query additional attributes from a data store, to specify issuance criteria, or to do 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.


Note that the **Target Session Mapping** configuration does not apply when the **No Mapping** option is chosen on the **Identity Mapping** screen.


11. On the **Protocol Settings** screen, click **Configure Protocol Settings**.


12. On the **OpenID Provider Info** screen, provide the scopes, the endpoints, and the authentication scheme; for example:



 **Note:** If you have chosen to load the metadata from the OP on the **General Info** screen, the **Scopes** field and all endpoints are pre-populated (provided that the metadata contains the information).


Field	Description
Scopes	<p>The scopes to be included in the authentication and token requests to the OP. Multiple space-separated values are allowed.</p> <p>The default value (without loading metadata from the OP) is <code>openid</code>.</p> <p> 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.</p>
Authorization Endpoint	<p>The authorization endpoint at the OP.</p> <p>You may enter a relative path (begin with a forward slash) if you have provided a base URL on the General Info screen.</p> <p>There is no default value (without loading metadata from the OP).</p>
OpenID Connect Login Type	<p>The OpenID Connect client profile of the client. This client represents PingFederate and is created and managed at the OP.</p> <ul style="list-style-type: none"> If the client is configured to support the Basic Client profile, select Code.

Field	Description
	<p>The resulting value of the <code>response_type</code> parameter is <code>code</code>.</p> <ul style="list-style-type: none"> If the client is configured to support the Implicit Client profile, select Form POST. <p>The resulting value of the <code>response_type</code> parameter is <code>id_token</code>.</p> <ul style="list-style-type: none"> If the client is configured to support the Implicit Client profile <i>and</i> the target application requires the associated access token, select Form POST with access token. <p>The resulting values of the <code>response_type</code> parameter are <code>id_token token</code>.</p> <p>The default selection (without loading metadata from the OP) is Code.</p>
Authentication Scheme	<p>The client authentication method that PingFederate uses. Applicable and visible only to clients supporting the Basic Client profile.</p> <ul style="list-style-type: none"> Select Basic to submit credentials via HTTP Basic authentication. Select POST to submit credentials via POST. Select Private Key JWT to authenticate via the <code>private_key_jwt</code> Client Authentication method, see Client Authentication in the OpenID Connect specification (openid.net/specs/openid-connect-core-1_0.html#ClientAuthentication). <p>The default selection (without loading metadata from the OP) is Basic.</p>
Authentication Signing Algorithm	<p>Select the algorithm that PingFederate uses to sign the JWT.</p> <p>Applicable and visible only when Private Key JWT is the chosen authentication scheme.</p> <p> Note: Based on the chosen signing algorithm, PingFederate selects the signing JSON Web Key (JWK) from its JWK Set (JWKS) at runtime.</p> <p>In order for the OP to validate the signed JWT, 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 <code><Base URL>/pf/JWKS</code>, where Base URL is defined on the Server Configuration > Server Settings > Federation Info screen. For example, if the Base URL field value is <code>https://www.example.com</code>, the PingFederate JWKS URL is <code>https://www.example.com/pf/JWKS</code>. You can pass the JWKS URL directly to the OP or have the OP contact the PingFederate OpenID Connect Metadata endpoint to obtain the information (see OpenID Provider configuration endpoint on page 476).</p>
Token Endpoint, User Info Endpoint, and JWKS URL	<p>Various OAuth 2.0 and OpenID Connect 1.0 endpoints at the OP. For more information, see openid.net/connect.</p> <p>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.)</p> <p>The User Info Endpoint field is optional. If omitted, PingFederate only has access to the end-user claims from the ID tokens.</p> <p>The JWKS URL is required in order for PingFederate to validated the inbound ID tokens from the OP.</p> <p>There are no default values (without loading metadata from the OP).</p>
Sign Request	<p>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.</p>

Field	Description
Request Signing Algorithm	<p>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 http://openid.net/specs/openid-connect-core-1_0.html#RequestObject.</p> <p>This check box is not selected by default, in which case PingFederate sends request parameters via multiple query parameters, unsigned.</p> <p>Select the algorithm that PingFederate uses to sign the request object.</p> <p>Applicable and visible only when the Sign Request check box is selected.</p> <p> Note: PingFederate automatically selects the signing JSON Web Key (JWK) based on the selected signing algorithm from its JWK Set (JWKS).</p> <p>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 <code><Base URL>/pf/JWKS</code>, where Base URL is defined on the Server Configuration > Server Settings > Federation Info screen. For example, if the Base URL field value is <code>https://www.example.com</code>, the PingFederate JWKS URL is <code>https://www.example.com/pf/JWKS</code>. You can pass the JWKS URL directly to the OP or have the OP contact the PingFederate OpenID Connect Metadata endpoint for it (see OpenID Provider configuration endpoint on page 476).</p>

13. Optional: Remain on the **OpenID Provider Info** screen and specify the request parameters that are allowed to be included in the authentication requests to the OP under **Request Parameters** (see [Configure request parameters and SSO URLs](#) on page 395).
14. Optional: On the **Overrides** screen, specify a default target URL and authentication context overrides.
15. On the **Activation & Summary** screen, review your connection settings.


When you finish setting up a connection, you may 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** screen for a new connection if you want to keep the configuration.

You can deactivate a connection at any time (for maintenance, for example). When a connection is inactive, all transactions to or from this partner are disabled.

In this use case, because PingFederate is essentially 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** screen. This registration should be associated with the client that represents PingFederate, the client that you have provided on the **General Info** screen. For more information, please refer to the documentation provided by the OP.

The **SSO Application Endpoint** near the top of the **Summary & Activation** screen is a sample URL that developers of the target application might use to invoke SSO for the connection. Additional query parameters might also be sent to the `/sp/startSSO.ping` endpoint.

 **Note:** If you have selected the **No Mapping** option on the **Identity Mapping** screen (because you are passing end-user claims to the target application through an authentication policy contract in an SP authentication policy), the **Summary & Activation** screen does not show the **SSO Application Endpoint** sample URL.

The target application can also invoke SSO requests by contacting the `/sp/init_login.ping` SP protocol endpoint.

For more information, see [Configure request parameters and SSO URLs](#) on page 395.

Configure request parameters and SSO URLs

On the **OpenID Provider Info** screen, administrators may define request parameters under **Request Parameters** for the following purposes:

- Allow custom request parameters to be include in the authentication requests to support 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. (Note that the OP may 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`).

Follow these steps to define one or more request parameters:

1. Add a request parameter under **Name**.
2. Define a default parameter value under **Value**.

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 `/sp/startSSO.ping` SP application endpoint. If the target application does not provide the parameter in its SSO URL (and no default value is specified), the parameter is not included in the authentication requests.

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.

If the target application does not provide the parameter in its 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 used.

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 that the `/sp/init_login.ping` 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 for more information.

4. Click **Add**.

Use the **Edit**, **Update**, and **Cancel** workflow to make or undo a change. Use the **Delete** and **Undelete** workflow to remove an entry or cancel the removal request.

5. Repeat these steps to define another request parameters.

Consider the following sample configuration:

REQUEST PARAMETERS			
Name	Value	Application Endpoint Override	Action
hd	example.org	<input type="checkbox"/>	Edit Delete
customMUEValue1	value one	<input type="checkbox"/>	Edit Delete
customMUEValue2	value two	<input type="checkbox"/>	Edit Delete
customOverrideOne	value can be overridden	<input checked="" type="checkbox"/>	Edit Delete
customOverrideTwo		<input checked="" type="checkbox"/>	Edit Delete
scope	address phone ext openid profile admin email	<input checked="" type="checkbox"/>	Edit Delete
acr_values	PasswordProtectedTransport	<input type="checkbox"/>	Edit Delete
prompt	login	<input type="checkbox"/>	Edit Delete
		<input type="checkbox"/>	Add

- 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 multi-valued 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** screen; for example:

```
https://sso.example.com/sp/startSSO.ping?PartnerIdpId=https%3A%2F%2Fsso.alpha.local%3A9031&customOverridableOne=foo
```

To construct a multi-valued request parameter, append the request parameter multiple times with different values; for example:

```
https://sso.example.com/sp/startSSO.ping?PartnerIdpId=https%3A%2F%2Fsso.alpha.local%3A9031&customOverridableOne=foo&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; therefore, 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 multi-valued 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 screen) 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 that while the target application can request different scopes, the OP may 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`.



Note: These examples use the `/sp/startSSO.ping` SP application endpoint. As needed, you may 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: When constructing SSO URLs, parameters with restricted characters must be URL-encoded.

For information about URL encoding, please refer to third party resources, such as [HTML URL-encoding Reference](#) (www.w3schools.com/tags/ref_urlencode.asp).

Query parameters vs. request object

By default, PingFederate sends all request parameters via multiple query parameters, unsigned. If the **Sign Request** check box is selected, PingFederate creates a signed JWT that contains claims representing the request parameters and passes the signed JWT as one query parameter, `request`, to the OP. Note that 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. (Note that the client authenticates via the HTTP Basic authentication scheme and initiates 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://sso.example.com/sp/
eyJpc3MiOiJodHRwczpcL1wvc3NvLmFscGhhLmxvY2FsOjkwMzEifQ/cb.openid
&state=e75n1lVU6Wa5TMmOwegDPSEI2iO9zd
&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 that the `client_id`, `response_type`, and `scope` request parameters are always passed to the OP as individual query parameter 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 http://openid.net/specs/openid-connect-core-1_0.html#RequestObject.

Configure IdP Auto-Connect

When your IdP partner is also using PingFederate 5 or higher (or is otherwise able to provide interoperable SAML 2.0 metadata via HTTP on demand), you may choose to use Auto-Connect for that partner (see [Using Auto-Connect](#)). This configuration can be shared among an unlimited number of SAML 2.0 partners.

 **Note:** You enable the SAML 2.0 Auto-Connect profile under System Settings (see [Choose roles and protocols](#) on page 151).

Once Auto-Connect is enabled on your PingFederate server, complete the configuration from the Main Menu under SP Configuration. This configuration entails:

- Setting up a common connection for all Auto-Connect partners
- Establishing a list of IdP partner domains authorized to use the connection

Initial setup for IdP Auto-Connect


The basic configuration for SP Auto-Connect requires only:

- Choosing a signing certificate for authentication requests and other SAML messages
- Configuring user-session creation information

All other partner-connection specifications are handled automatically at runtime.

Choose a certificate for IdP Auto-Connect

For Auto-Connect runtime processing, authentication requests and SLO messages must be signed, because they are sent over either the POST or redirect bindings (see [SAML 2.0 profiles](#) on page 34 in the “Supported Standards” chapter of the *Get started with PingFederate* guide).

 **Note:** The signing certificate is embedded in your server's Auto-Connect metadata (see [Using Auto-Connect](#) on page 79); there is no need to exchange certificates with your partners.

You can use the same certificate used for signing metadata (see [Configure metadata signing](#) on page 156). If you use a different certificate, ensure that it meets Auto-Connect validation requirements (see [Auto-Connect security model](#) on page 80).

To specify a certificate:

1. Select a signing certificate from the list.

If you have not yet created or imported your certificate into PingFederate, click **Manage Certificates** (see [Manage digital signing certificates and decryption keys](#) on page 184).

2. 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.

Configure user-session creation for IdP Auto-Connect

Configuring user-session creation for Auto-Connect is similar to configuring the same settings for regular partner connections.

- Click **Configure User-Session Creation** to continue.

For configuration information, refer to sections under *Configure user-session creation* on page 347.



Note: Attributes sent from the IdP via Auto-Connect are passed to your applications, regardless of whether they are listed in the attribute contract (see *Define an attribute contract* on page 348).

Review IdP Auto-Connect settings

When you finish configuring your IdP Auto-Connect initial setup, you may choose to activate the common connection immediately on the Activation & Summary screen. (No runtime processing occurs until your partner's Auto-Connect gateway is also established and a user initiates an SSO or SLO event.)



Important: Regardless of whether you choose to activate a newly configured connection now or later, you must click **Save** on the Activation & Summary screen if you want to keep the configuration.

You can deactivate the connection at any time (for maintenance, for example). While a connection is inactive, all SSO or SLO transactions to or from Auto-Connect partners are disabled.

To change a Connection Status:

- Select Active or Inactive and then click **Save**.

To modify a setting:

1. Locate the currently configured setting under one of the summary headings and then click the subheading above the data.



Note: Changes made to Auto-Connect settings will be out of sync, temporarily, with metadata caches that any currently active partners might be using. If your connection is in production, you might wish to lower your server's metadata lifetime in advance of making configuration changes (see *Configure metadata lifetime* on page 157).

2. Change the information and click **Save**, if available.

If **Save** is not available, additional, dependent changes are required; click **Next** or **Done** until you reach a screen containing a **Save** button. Then click **Save** and continue as needed until you return to the Main Menu.

Specify allowed IdP domains

This screen provides PingFederate® with a list of trusted domain names of your Auto-Connect partners.

Normally, when PingFederate receives an SSO request from a web application at your site (see */sp/startSSO.ping*), the runtime engine completes the connection automatically using metadata obtained from a standard, public location—`http://saml.<domain_name>`. (See *Using Auto-Connect*.) Alternatively, if an Auto-Connect partner elects not to use the standard location, you can supply the applicable URL.

To add a domain:

- Enter a Domain Name and click **Add**.

To specify a URL for metadata retrieval:

1. Click **Advanced View**.
2. Enter the Domain Name if you have not already done so.
3. Enter the Metadata Service URL.

This entry must be obtained from your Auto-Connect partner.

4. Click **Add**.



Note: Once you have added the URL, you cannot return to the **Basic View** unless you first remove the URL value using the procedure below.

To edit an entry:

1. Click **Edit** under Action for the entry.
2. Make your change and click **Update**.

To delete an entry:

- Click **Delete** under Action for the entry.

WS-Trust STS configuration

The PingFederate® WS-Trust Security Token Service (STS) provides security-token validation and creation to extend SSO access to identity-enabled Web Services (see [Web Services standards](#) on page 46).

The chapter provides instructions for configuring the WS-Trust STS.

- [Server settings](#) on page 400
- [IdP configuration for STS](#) on page 401
- [SP configuration for STS](#) on page 412

Server settings

To use the PingFederate WS-Trust STS for partner connections, start by enabling the WS-Trust protocol under Server Settings on the Roles and Protocols screen (see [Choose roles and protocols](#) on page 151). Once the protocol is enabled, 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 corresponding browser-based SSO protocols).

In addition, also under Server Settings, you have the option of requiring authentication globally for access to STS endpoints—(see [Configure STS authentication](#) on page 400).

Enable the WS-Trust protocol

1. Go to the **Server Configuration > Server Settings > Roles & Protocols** screen.
2. If you have not done so, enable the applicable federation role (or roles) for your deployment.
3. Under the role (or roles) that you require STS processing, select the **WS-Trust** check box to enable the protocol.



Note: PingFederate supports the STS with or without selections of other browser-based SSO protocols. The handling of SAML 1.1 and 2.0 tokens is independent of the supported browser-based SSO protocols shown on the same screen.

4. Enter your SAML federation IDs on the **Server Configuration > Server Settings > Federation Info** screen (unless these IDs are already established for 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 as needed.

Configure STS authentication

On the **Server Configuration > Server Settings > WS-Trust STS Settings** screen, you may configure PingFederate® to require that client applications provide credentials to access the STS.

While this is an optional configuration, it is recommended for 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 may 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 SP partner connections, or token-to-token mappings, using more fine grained controls, at the connection level via Issuance Criteria.

1. On the **WS-Trust STS Settings** screen, click **Configure WS-Trust STS Authentication**.
2. On the **Authentication Methods** screen, select HTTP Basic, mutual SSL/TLS authentication, or both.

If both HTTP Basic and mutual SSL/TLS authentication are selected, *all* clients must provide credentials for both mechanisms.



Important: If you choose mutual SSL/TLS authentication, you must configure a secondary PingFederate HTTPS port (pf.secondary.https.port) in the `run.properties` file (see [Configure PingFederate properties](#) on page 93).

3. If you have chosen HTTP Basic, manage user accounts on the **HTTP Basic Authentication** screen.

Click **Create User** and enter a username and its password on the **User Account** screen. Repeat to create additional user accounts for your client applications.

On the **HTTP Basic Authentication** screen, you can also delete user accounts and update their passwords.

4. If you have chosen mutual SSL/TLS authentication, click **Configure Mutual SSL Authentication** on the **Mutual SSL Authentication** screen.

- a) On the **Authentication Options** screen, select to restrict access by the subject DN or issuer of the client certificate.

If both options are selected, the client certificate used for authentication to the STS endpoints must meet both sets of restrictions.

- b) If you have chosen to restrict by the subject DN, enter one or more subject DNs on the **Allowed Subject DNs** screen.

On the **Allowed Subject DNs** screen, you may edit or delete existing entries but you must keep at least one subject DN.

- c) If you have chosen to restrict by the certificate issuer, select one or more client certificate on the **Allowed Issuer Certificates** screen.

If you have not yet imported the client certificate, click **Manage Certificates** to do so.

On the **Allowed Issuer Certificates** screen, you may remove existing entries but you must keep at least one issuer.

- d) On the **Summary** screen, review your mutual SSL/TLS authentication settings and then click **Done**.

5. When you finish configuring WS-Trust STS settings, you can review the configuration on its **Summary** screen.

If you want to keep your changes, click **Save**.

IdP configuration for STS

This section covers the IdP configuration for the PingFederate WS-Trust STS, which involves:

- [Manage token processors](#) on page 401
- [Manage STS request parameters](#) on page 405 (Optional)
- [Configure SP connections for STS](#) on page 405

Manage token processors

The PingFederate® Security Token Service (STS) uses token processors to validate incoming tokens and token requests. 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 46.)

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 also deploy additional token translators from [Ping Identity website](http://www.pingidentity.com/en/products/downloads.html) (www.pingidentity.com/en/products/downloads.html).

You manage token processor instances on the **IdP Configuration > Token Processors** screen.

- 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**.



Note: Automatic multi-connection error checking occurs by default whenever you access this screen. The intent is to verify that configured connections have not been adversely affected by changes made here.

If you experience noticeable delays in accessing this screen, you can optionally disable automatic connection validation on the **Server Configuration > Server Settings > System Options** screen.

For simplicity, this topic focuses on configuring an instance of one of the integrated token processors. For add-on token generators, please refer to the online documentation referenced in the download package.

Select a token processor type

The first step in creating a token-processor instance is choosing the processor type.

- On the **Type** screen, configure the basics of this token processor instance.
 - a) Enter a name and ID for this token processor instance.
 - b) Select the token-processor type from the list.
 - c) 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 screens.

Configure a token processor instance

Depending on the selected token processor, the **Instance Configuration** screens presents you with different parameters.

For token processors bundled with PingFederate[®], refer to one of the following sections:


- For token processors bundled with PingFederate[®], refer to one of the following sections:
 - [Configure a Username Token Processor instance](#) on page 402
 - [Configure a Kerberos Token Processor instance](#) on page 403
 - [Configure an OAuth Token Processor instance](#) on page 403
 - [Configure a JSON Web Token Processor instance](#) on page 403
 - [Configure a SAML Token Processor instance](#) on page 404
- For add-on processors, please refer to the online documentation referenced in the download package or the [Knowledge Center](#).

Configure a Username Token Processor instance

The integrated Username Token Processor accepts and validates username security tokens.

- On the **Instance Configuration** screen, configure the basics of this token processor instance.
 - a) Click **Add a new row to 'Credential Validators'** to define a credential-authentication mechanism instance for the token processor.
 - b) Select a Password Credential Validator instance from the list and click **Update**.

Add as many validators as necessary. Click **Move up** and **Move down** 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 is able to 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 lockout those accounts in their source locations.


If you have not yet defined the desired Password Credential Validator instance, click **Manage Password Credential Validators** to do so.

Configure a Kerberos Token Processor instance

The integrated Kerberos Token Processor accepts and validates Kerberos tokens via a configured Kerberos realm. Moreover, it supports authentication mechanism assurance from Active Directory domain service, thus making it possible to restrict access to users authenticating through specific mechanisms. For more information, see [Authentication mechanism assurance](#) on page 427.

- On the **Instance Configuration** screen, select the applicable domain from the **Domain/Realm Name** list.

An Active Directory 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 [Configure Active Directory domains or Kerberos realms](#) on page 203.)

 **Note:** Kerberos tickets can be accepted from domains other than the domain configured in the Token Processor, provided that there is a transient, two-way trust. This trust exists by default when domains are joined within a single server forest (see [Multiple-domain support](#) on page 203).

Configure an OAuth Token Processor instance

The PingFederate® STS provides validation for OAuth 2.0 bearer tokens. Generally, a client would send a base64-encoded access token in order to receive a SAML token in exchange. To use this token processor, you must first configure an Access Token Management (ATM) instance.

(For more information about PingFederate OAuth authorization server and access token management, see [About OAuth](#) on page 61 and [OAuth access token management](#) on page 245.)

- On the **Instance Configuration** screen, configure the basics of this token processor instance.
 - a) Select an Access Token Management instance from the list.

The token processor instance uses the selected ATM instance to validate the OAuth bearer access tokens.

- b) Optional: Select the **Scope Value as Single String** check box.

If selected, the scope value is returned as a single space-delimited set of string values; otherwise, scope values are returned as a multi-valued attribute.

Configure a JSON Web Token Processor instance

The PingFederate® STS provides validation for JSON web tokens (JWTs).

- On the **Instance Configuration** screen, provide the required information.

Refer to the following table for information about each field.

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.
Expiry Tolerance	The amount of time (in seconds) to allow for clock skew between servers. Valid range is 0 to 3600.

Configure 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® STS validates digital signatures using all trusted certificate authorities (CAs) imported into PingFederate. As needed, you may restrict the signature verification process by subject DNs or issuers (or both) to limit the token requests accepted for this token processor instance.

In addition, you must indicate a unique identifier for the PingFederate STS. Once defined, incoming SAML tokens must contain this ID in its audience element in order for them to be accepted by this token processor instance.

- On the **Instance Configuration** screen, configure the basics of this token processor instance.
 - a) Enter the URI that uniquely identifies your federation gateway for this SAML protocol in the **Audience** field.

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'** and then enter one or more issuers.


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'** and then enter one or more subject DNs.

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.
-  **Important:** 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.

Extend a token processor contract

Token processors allow administrators to add to a built-in list of user attributes that the processor returns from an incoming token; it is essentially an extended processor attribute contract. Note that the **Extended Contract** screen shows a different list of attributes under **Core Contract**, depending on the token processor selected.

- Enter the name of the desired attribute and 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 this step as needed to add another attribute.

Set attribute masking

On the **Token Attributes** screen, you can choose to mask attribute values that PingFederate® logs from this processor instance at runtime.

- Optional: Under **Mask Log Values**, select the attribute whose value you want to mask in logs.

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 related check box under the list of attributes.

Review the token processor configuration

1. On the **Summary** screen, review your configuration, modify as needed, and click **Done** to exit the **Create Token Processor Instance** task.
2. On the **Manage Token Processor Instances** screen, click **Save** to retain the configuration of the token processor instance.

If you want to exit without saving the configuration, click **Cancel**.

Manage STS request parameters

As an option for configuring PingFederate® to act as a WS-Trust STS, you can define sets of RST metadata parameters that can be used to map attribute values into issued security tokens. After these *request contracts* are defined, you can make them available when configuring WS-Trust STS settings in the SP connections (see [Select a request contract](#) on page 407).

You manage request contract on the **IdP Configuration > STS Request Parameters** screen.

- To configure a new set of request parameters, click **Add New Request Contract**.
- To modify an existing request contract, select it by its name under **Contract Name**.
- To review the usage of an existing request contract, click **Check Usage** under **Action**.
- To remove an existing request contract or to cancel the removal request, click **Delete** or **Undelete** under **Action**.

Create a request contract

On the **Create Request Contract** screen, identify the contract and define parameters that will be available in token requests (as associated with this contract for partner connections).

1. Enter the contract name and contract ID.

Once the request is saved, the name and ID cannot be modified.

2. Enter the parameter that will be included in RSTs and click **Add**.

You must add at least one parameter. Repeat this step to add more parameters.

Once the request is saved, you may add, modify, or remove parameters but you must keep at least one parameter.

Configure SP connections for STS

You can configure an STS connection to an SP partner either in conjunction with browser-based SSO or independently.

- To configure an STS connection, select the **WS-Trust STS** check box.

The **WS-Trust STS** option is only available after you enable the **WS-Trust** role on the **Server Configuration > Server Settings > Roles & Protocols** screen.

Configure protocol settings for IdP STS

1. Enter a URL for your partner's Web Service in the text field and then click **Add**.

This identifier is compared to the *AppliesTo* element in the Requests for Security Token (RST) messages and may be either a complete URL or a base URL for matching variable ports or paths.

Repeat this step to add additional identifiers as needed.

2. Select any of the following options that are applicable to your use case.

Option	Description
OAuth Assertion Profiles	<p>When selected, four additional token-type requests become available based on these OAuth grant types:</p> <ul style="list-style-type: none"> • 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 <p>See STS OAuth integration on page 60 for more information on the use of these token-type requests.</p>
Default Token Type	<p>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:</p> <ul style="list-style-type: none"> • SAML 2.0

Option	Description
	<ul style="list-style-type: none"> • SAML 1.1 • SAML 1.1 for Office 365 <p>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).</p>
Generate Key for SAML Holder of Key Subject Confirmation Method	<p>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.</p> <p>For information about HoK assertions, see Web Services Security SAML Token Profile (docs.oasis-open.org/wss-m/wss/v1.1.1/os/wss-SAMLTOKENProfile-v1.1.1-os.html).</p> <p>This option does not apply to OAuth assertion profiles.</p>
Encrypt SAML 2.0 Assertion	<p>When selected, the STS encrypts the SAML 2.0 assertion. Applicable only to SAML 2.0 security token.</p> <p>This option does not apply to OAuth assertion profiles.</p>

Set a token lifetime

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.

- Optional: Override the default values for the following fields:

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.

Configure token creation

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 (the "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 SSO, the mapping can be augmented using local data stores, variable or constant text, or expressions.

Details of this configuration are handled under the **Token Creation** task.

- To continue, click **Configure Token Creation**.

Define an attribute contract for IdP STS

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 (see [Attribute contracts](#) on page 71). You identify these attributes on the **Attribute Contract** screen.

1. Enter the attribute name in the text box.

Attribute names are case-sensitive and must correspond to the attribute names (including claims) expected by the requesting WSC.



Tip: The Format attribute associated with the NameID element in outgoing SAML tokens may be set when needed by adding an attribute called SAML_NAME_FORMAT. The value of that attribute can then be mapped later (see [Fulfill the attribute contract for token creation](#) on page 409).

For information about the NameID elements and applicable URI values, locate the SAML 2.0 specification at www.oasis-open.org/standards.



Tip: 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 IdP application—authentication context. Map a value for the authentication context on the attribute-mapping screen later in the configuration, from any available attribute source, including the RST if a requested context is specified as a request parameter (see [Fulfill the attribute contract for token creation](#) on page 409).

- 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** on the **WS-Trust STS > Protocol Settings** screen.

Change the default namespace selection if you and your SP partner have agreed to a specific namespace.



Note: As needed, 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. (Restart of PingFederate is required to activate any changes made to this file.)

(For more information about attribute namespace, see [Attribute contracts](#) on page 71.)

- Click **Add**.
- Repeat until all applicable attributes are defined.

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.

Select a request contract

This optional setting on the **Request Contract** screen allows you to use XML parameters contained in RSTs for token-attribute mapping (see [Manage STS request parameters](#) on page 405).

- If you are not using request parameters, click **Next** to continue.
- To use request parameters, select the check box and choose a request contract from list.

If the contract you want is not shown, click **Manage STS Request Parameters**.

When selected, you may choose the request contract as the attribute source in the **IdP Token Processor Mapping** configuration later in the setup.

If you are editing an existing configuration, you may change the request contract or disable this optional setting. These changes may require additional configuration changes in subsequent tasks.

Manage IdP token processor mappings

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.

You can map one or more token processor instances into an SP connection to satisfy multiple session-management requirements where needed. (The same token processor instances may also be mapped in multiple SP connections.)

When token processor instances are restricted to certain virtual server IDs, the allowed IDs are displayed under **Virtual Server IDs**.

- To map a token processor instance, click **Map New Token Processor Instance**.
- 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**.
- 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**.

Select a token processor instance


On the **Token Processor Instance** screen, choose an instance of a deployed token processor that suits your requirements for this connection.

1. Select a token processor instance from the list.

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.


When selected, the administrative console adds a new set of sub tasks (the **Override Instance** screen and its sub tasks).

 **Tip:** 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 plug-in configurations](#) on page 69.

If you are editing a currently mapped token processor 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.


Override a token processor instance

On the **Override Instance** screen, you start a series of sub tasks to 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.

- Click **Override Instance Settings**.

On each of the settings screens, select the **Override** check box, make your changes, and then click **Next**. When you are finished, click **Done** to continue with the rest of the mapping configuration.

 **Note:** The override setting screens are functionally identical to those used for creating a new token processor instance (see [Manage token processors](#) on page 401).

If you are editing a currently mapped token processor instance, click **Override Instance Settings** to reconfigure any overridden settings for this connection. You may also remove all overridden settings on a per-screen basis by clearing the **Override** check box near the top of the screen.

Restrict a token processor to certain virtual server IDs

When you multiplex one connection for multiple environments (see [Multiple virtual server IDs](#) on page 84), you can enforce authentication requirements by restricting a token processor to certain virtual server IDs on the **Virtual Server IDs** screen. By default, no restriction is imposed.

1. Select the **Restrict Virtual Server IDs** check box.
2. Select one or more virtual server IDs that you want to allow for this token processor.

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.

Select 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®.

The values supplied by the token processor are shown under **Token Processor Contract** on the **Attribute Retrieval** screen.

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.

- Select the **Retrieve additional attributes from data stores ...** 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 ...** option if you do *not* require data store query.

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.

Configure attribute sources and user lookup for token creation

Attribute sources are specific data store or directory locations containing information that may be needed for the attribute contract. 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.

- Click **Add Attribute Source** and then follow a series of sub tasks to complete the configuration.

For step-by-step instructions, see [Choose a data store](#) on page 508.

Repeat to add additional attribute sources.

If you are editing a currently mapped token processor instance, you can add, remove, or reorder attribute sources, which may require additional configuration changes in subsequent tasks.

Fulfill the attribute contract for token creation

On the **Attribute Contract Fulfillment** screen, map values to the attributes defined for the contract. These are the values that will be included in the SAML security tokens sent to the SP.

- For each attribute, select a source from the list and then choose or enter a value.


- **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.

 **Note:** 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 (see **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 **Server Configuration > Server Settings > WS-Trust STS Settings** screen.

- **Request**

When selected, the **Value** list is populated 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** screen. 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 Custom**

When selected, the **Value** list is populated 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 SAML security token.

- **Expression** (when enabled)

This option provides more complex mapping capabilities; for example, 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 (see **Text**).

Note that expressions are not enabled by default. For more information about enabling and editing OGNL expressions, see [Attribute mapping expressions](#) on page 497.

- **Text**

When selected, the text you enter is mapped to the value of the attribute in the SSO 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 this 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.

All attributes must be mapped.


If you are editing a currently mapped token processor instance, you can update the mapping configuration, which may require additional configuration changes in subsequent tasks.

Specify issuance criteria for token creation

On the **Issuance Criteria** screen, you may define the criteria that PingFederate® can evaluate to determine whether to issue a SAML security token for a user. This optional token authorization configuration can be used to restrict who can access identity-enabled Web Services (For more information, see [About token authorization](#) on page 74.)

Issuance criteria and OGNL expressions


When defining multiple criteria using the user-interface elements, all criteria must pass in order for a user to be authorized. Source attribute values are evaluated against specified values using an exact match algorithm. The **multi-value contains** (or **multi-value does not contain**) conditions (intended for attributes that may contain multiple values) return true if the source attribute contains the exact specified value. For more complex evaluations, including conditional criteria and partial matching, define criteria using OGNL expressions.

 **Important:** When you multiplex one connection for multiple environments (see [Multiple virtual server IDs](#) on page 84), consider using an OGNL expression to verify the virtual server ID in conjunction with other conditions, such as group membership information, to protect against unauthorized access (see [Sample OGNL expressions](#) on page 498).

Note that expressions are not enabled by default. For more information about enabling and editing OGNL expressions, see [Attribute mapping expressions](#) on page 497.

1. Select the source of the attribute that you want PingFederate to evaluate from the **Source** list.

When you select a source, the **Attribute Name** list is populated with the associated attributes or properties, as described in the following table.

Source	Attribute name
Token	Populated with the attributes from the token processor instance.
Context	Populated with the available context of the transaction.  Note: 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.
Request (if a request contract is configured)	Populated with parameters from the request contract.
LDAP, JDBC, or Custom (if a data store is configured)	Populated with attributes that you have selected in the attribute source configuration.
Mapped Attributes	Populated with attributes from the attribute contract.

2. Select an attribute from the **Attribute Name** list.



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 **Server Configuration > Server Settings > WS-Trust STS Settings** screen.

3. Select the evaluation method from the **Condition** 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



Note: Use one of the *multi-value* conditions when you want PingFederate to verify that an attribute contains or does not contain the specified value in its attribute list. Using a single-value condition when an attribute has multiple values causes the condition to evaluate consistently to false.



Important: When using the **STS SSL Client Certificate's Subject DN** attribute, you must select one of the following conditions: **equal to DN**, **not equal to DN**, **multi-value contains DN**, or **multi-value does not contain DN**. These operators normalize the DN before comparison to accommodate for different string representations that are still considered equivalent (for example, case sensitivity, or whitespace).

4. Enter the exact value in the **Value** field.

Wildcards (for partial matching) are not allowed.



Tip: You can use OGNL expressions to perform more complex evaluations, including partial matching.

5. Optional: Use the **Error Result** field to customize an error code or error message for use if the token authorization fails.

Errors result in a SOAP message returned. 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* (www.w3.org/TR/2000/NOTE-SOAP-20000508/#_Toc478383507).



Tip: 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.

If you leave this field blank, a default `ACCESS_DENIED` error result is used if the authorization fails.

6. Click **Add**.

7. Repeat these steps if you want to add additional criteria.



Important: All criteria must pass in order for a user to be authorized.

8. Optional: Click **Show Advanced Criteria** to configure more complex evaluations using OGNL expressions.

- a) Enter an OGNL expression in the **Expression** field.
- b) Enter an error code or an error message in the **Error Result** field.



Note: If the OGNL expression resolves to a string value (rather than `true` or `false`), the returned value overrides the **Error Result** field value.

- c) Click **Add**.
- d) Click **Test**, enter values in the applicable fields, and test the resulting output for the expression.

e) Repeat these steps if you want to add additional criteria using expressions.



Important: All criteria must pass in order for a user to be authorized.

If you want conditional criteria, add only one criterion with such processing logic.

If you are editing a currently mapped token processor instance, you can add, remove, or edit any issuance criteria.

Review the IdP token processor mapping

When you have finished adding a new or modifying an existing token processor instance, you can review the configuration on its **Summary** screen.

- If you need to make any changes, click the heading over the information you want to edit.
- If you are creating a new connection and want to keep your configuration, click **Done**.
- If you are editing an existing configuration and want to keep your changes, click **Save**.

Select 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 RST messages, you can choose whether to continue or abort the token-creation process.

Note that the **Error Handling** screen is presented only if a request contract is configured on the **Request Contract** screen for this connection.

Review the token creation configuration

When you are finished with the **Token Creation** task, you can review the configuration on its **Summary** screen.

- If you need to make any changes, click the heading over the information you want to edit.
- If you are creating a new connection and want to keep your configuration, click **Done**.
- If you are editing an existing configuration and want to keep your changes, click **Save**.

Review the IdP STS settings

When you are finished with the **WS-Trust STS** task, you can review the configuration on its **Summary** screen.

- If you need to make any changes, click the heading over the information you want to edit.
- If you are creating a new connection and want to keep your configuration, click **Done**.
- If you are editing an existing configuration and want to keep your changes, click **Save**.

SP configuration for STS

This section covers the SP configuration for STS, including:

- [Manage token generators](#) on page 412
- [Configure IdP connections for STS](#) on page 414

Manage 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.

(For more information about WS-Trust, see [Web Services standards](#) on page 46.)

PingFederate comes bundled with the SAML 1.1 Token Generator and SAML 2.0 Token Generator.

You can also deploy additional token translators from [Ping Identity website](#) (www.pingidentity.com/en/products/downloads.html).

You manage token generator instances on the **SP Configuration > Token Generators** screen.

- 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**.



Note: Automatic multi-connection error checking occurs by default whenever you access this screen. The intent is to verify that configured connections have not been adversely affected by changes made here.

If you experience noticeable delays in accessing this screen, you can optionally disable automatic connection validation on the **Server Configuration > Server Settings > System Options** screen.

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.

Select a token generator type

The first step in creating a token-generator instance is choosing the generator type.

- On the **Type** screen, configure the basics of this token generator instance.
 - a) Enter a name and ID for this token generator instance.
 - b) Select the token-generator type from the list.
 - c) 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 screens.

Configure a token generator instance

Depending on the selected token generator, the **Instance Configuration** screens presents you with different parameters.

For the integrated SAML 1.0 and 2.0 Token Generators, refer to the following table and specify parameters for generated SAML tokens.

(For add-on processors, please refer to the online documentation referenced in the download package or the [Knowledge Center](#).)

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 on the Server Configuration > Serer Settings > Roles & Protocols screen.
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 (see Manage digital signing certificates and decryption keys on page 184).
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 instead.
Include Raw Key in KeyValue	If selected, the raw key is included in the KeyInfo element as well.

Field	Instructions
Audience	This is a unique identifier for the target Web Service, used for the audience element of the generated SAML token.
Confirmation Method	Choose from among available methods: <ul style="list-style-type: none"> • urn...cm:sender-vouches (default) • urn...cm:bearer • urn...cm:holder-of-key For more information, see WSS SAML Token Profile (docs.oasis-open.org/wss-m/wss/v1.1.1/os/wss-SAMLTokenProfile-v1.1.1-os.html).
Encryption Certificate	The Web Service Provider's public certificate for encryption is required <i>only</i> 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 (see Manage certificates from partners on page 190).

Extend 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.

Token generators allow administrators to add to a built-in list of user attributes that the generator includes in the outgoing token; it is essentially an extended generator-attribute contract. Note that the **Extended Contract** screen shows a different list of attributes under **Core Contract**, depending on the token generator selected.

- Enter the name of the desired attribute and click **Add**.
Repeat this step as needed to add another attribute.

Review the token generator configuration

1. On the **Summary** screen, review your configuration, modify as needed, and click **Done** to exit the **Create Token Generator Instance** task.
2. On the **Manage Token Generator Instances** screen, click **Save** to retain the configuration of the token generator instance.

If you want to exit without saving the configuration, click **Cancel**.

Configure IdP connections for STS

You can configure an STS connection to an IdP partner either in conjunction with browser-based SSO or independently.

- To configure an STS connection, select the **WS-Trust STS** check box.



Note: Before this option can be selected, the WS-Trust protocol must be enabled in Server Settings (see [Server settings](#) on page 400).

Configure protocol settings for SP STS

- On the **Protocol Settings** screen, choose whether to only validate incoming SAML tokens or to validate and then also generate local tokens to enable SSO access to Web Services at your site.

If you choose to generate local tokens as well, you must set up at least one token generator.

Configure token generation

For the PingFederate[®] STS to issue a security token that meets identity requirements of Web Services at your site, you must indicate what user attributes are included in the incoming token. As with Browser SSO, the mapping can be augmented using local data stores, variable or constant text, or expressions.

Details of this configuration are handled under the **Token Generation** task.

- To continue, click **Configure Token Creation**.

Define an attribute contract for SP STS

An attribute contract is the set of user attributes expected in incoming SAML assertions (see [Attribute contracts](#) on page 71). You identify these attributes on the **Attribute Contract** screen.

Optionally, you can mask the values of attributes (other than `SAML_SUBJECT`) in logs that PingFederate writes when it receives security tokens.

1. Enter the attribute name in the text box.

Attribute names are case-sensitive and must correspond to the attribute names expected by the requester.

2. Optional: Select the check box under **Mask Values in Log**.

3. Click **Add**.

4. Repeat until all applicable attributes are defined.

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.

Manage 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.

As needed, you can map one or more token generator instances into an IdP connection to satisfy different token requirements by the Web Services at your site. (The same token generator instances may also 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**.

- To map a token generator instance, click **Map New Token Generator Instance**.
- To edit the mapping configuration of a token generator instance, open it by clicking on its name, select the setting that you want to reconfigure, and complete the change.
- To remove a token generator instance 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 configuration, click **Done**.
- If you are editing an existing configuration and want to keep your changes, click **Save**.

Select a token generator instance


On the **Token Generator Instance** screen, choose an instance of a deployed token processor that suits your requirements for this connection.

1. Select a token generator instance from the list.

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.

2. Select the **Override Instance Settings** check box if you want to customize one or more token processor settings for this connection alone.

When selected, the administrative console adds a new set of sub tasks (the **Override Instance** screen and its sub tasks).

 **Tip:** 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 plug-in configurations](#) on page 69.

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.

Override a token generator instance

On the **Override Instance** screen, you start a series of sub tasks to override token generator settings specifically for this 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.

- Click **Override Instance Settings**.

On each of the settings screens, select the **Override** check box, make your changes, and then click **Next**. When you are finished, click **Done** to continue with the rest of the mapping configuration.



Note: The override setting screens are functionally identical to those used for creating a new token generator instance (see *Manage token generators* on page 412).

Restrict a token generator to certain virtual server IDs

When you multiplex one connection for multiple environments (see *Multiple virtual server IDs* on page 84), you can enforce integration requirements by restricting a token generator to certain virtual server IDs on the **Virtual Server IDs** screen. By default, no restriction is imposed.

1. Select the **Restrict Virtual Server IDs** check box.
2. Select one or more virtual server IDs that you want to allow for this token generator.

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.

Select an attribute retrieval method for token generation

For token generation, you can query 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 under **Attribute Contract** on the **Attribute Retrieval** screen.

- 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 *Choose a data store* on page 508.

If you are editing a currently mapped token generator instance, you can change the mapping method, which may require additional configuration changes in subsequent tasks.

Fulfill the attribute contract for token generation

On the **Token Generator Contract Fulfillment** screen, map values to the attributes defined for the contract. These are the values that the Web Services require.

- For each attribute, select a source from the list and then choose or enter a value.


- **Assertion**


When selected, the **Value** list is populated 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 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 local token.

 **Note:** 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 (see **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 **Server Configuration > Server Settings > WS-Trust STS Settings** screen.

- **LDAP, JDBC, or Custom**

When selected, the **Value** list is populated 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; for example, 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 (see **Text**).

Note that expressions are not enabled by default. For more information about enabling and editing OGNL expressions, see [Attribute mapping expressions](#) on page 497.

- **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 data store, when applicable, using this syntax:

```
${ds.attribute}
```

where *attribute* is any of the attributes that you have selected from the data store.

All attributes must be mapped.


If you are editing a currently mapped token generator instance, you can update the mapping configuration, which may require additional configuration changes in subsequent tasks.

Specify issuance criteria for token generation

On the **Issuance Criteria** screen, you may define the criteria that PingFederate® can evaluate to determine whether to issue a local security token (see [About token authorization](#) on page 74). This optional token authorization configuration can be used to restrict who can access identity-enabled Web Services.

Issuance criteria and OGNL expressions


When defining multiple criteria using the user-interface elements, all criteria must pass in order for a user to be authorized. Source attribute values are evaluated against specified values using an exact match algorithm. The **multi-value contains** (or **multi-value does not contain**) conditions (intended for attributes that may contain multiple values) return true if the source attribute contains the exact specified value. For more complex evaluations, including conditional criteria and partial matching, define criteria using OGNL expressions.

 **Important:** When you multiplex one connection for multiple environments (see [Multiple virtual server IDs](#) on page 84), consider using an OGNL expression to verify the virtual server ID in conjunction with other conditions, such as group membership information, to protect against unauthorized access (see [Sample OGNL expressions](#) on page 498).


Note that expressions are not enabled by default. For more information about enabling and editing OGNL expressions, see [Attribute mapping expressions](#) on page 497.

1. Select the source of the attribute that you want PingFederate to evaluate from the **Source** list.

When you select a source, the **Attribute Name** list is populated with the associated attributes or properties, as described in the following table.

Source	Attribute name
Assertion	Populated with the attributes from the incoming SAML token.
Context	Populated with the available context of the transaction.
	 Note: 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.
LDAP, JDBC, or Custom (if a data store is configured)	Populated with attributes that you have selected from the data store.
Mapped Attributes	Populated with attributes from the attribute contract.


2. Select an attribute from the **Attribute Name** list.


 **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 **Server Configuration > Server Settings > WS-Trust STS Settings** screen.

3. Select the evaluation method from the **Condition** 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

 **Note:** Use one of the *multi-value* conditions when you want PingFederate to verify that an attribute contains or does not contain the specified value in its attribute list. Using a single-value condition when an attribute has multiple values causes the condition to evaluate consistently to false.

 **Important:** When using the **STS SSL Client Certificate's Subject DN** attribute, you must select one of the following conditions: **equal to DN**, **not equal to DN**, **multi-value contains DN**, or **multi-value does not contain DN**. These operators normalize the DN before comparison to accommodate for different string representations that are still considered equivalent (for example, case sensitivity, or whitespace).


4. Enter the exact value in the **Value** field.

Wildcards (for partial matching) are not allowed.

 **Tip:** You can use OGNL expressions to perform more complex evaluations, including partial matching.

5. Optional: Use the **Error Result** field to customize an error code or error message for use if the token authorization fails.


Errors result in a SOAP message returned. 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* (www.w3.org/TR/2000/NOTE-SOAP-20000508/#_Toc478383507).

-  **Tip:** 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.

If you leave this field blank, a default `ACCESS_DENIED` error result is used if the authorization fails.


6. Click **Add**.

7. Repeat these steps if you want to add additional criteria.


-  **Important:** All criteria must pass in order for a user to be authorized.

8. Optional: Click **Show Advanced Criteria** to configure more complex evaluations using OGNL expressions.

- a) Enter an OGNL expression in the **Expression** field.
- b) Enter an error code or an error message in the **Error Result** field.

 **Note:** If the OGNL expression resolves to a string value (rather than `true` or `false`), the returned value overrides the **Error Result** field value.

- c) Click **Add**.
- d) Click **Test**, enter values in the applicable fields, and test the resulting output for the expression.
- e) Repeat these steps if you want to add additional criteria using expressions.

 **Important:** All criteria must pass in order for a user to be authorized.

If you want conditional criteria, add only one criterion with such processing logic.

Review the SP token generator mapping

When you have finished adding a new or modifying an existing token generator instance, you can review the configuration on its **Summary** screen.

- If you need to make any changes, click the heading over the information you want to edit.
- If you are creating a new connection and want to keep your configuration, click **Done**.
- If you are editing an existing configuration and want to keep your changes, click **Save**.

Review the token generation configuration

When you are finished with the **Token Generation** task, you can review the configuration on its **Summary** screen.

- If you need to make any changes, click the heading over the information you want to edit.
- If you are creating a new connection and want to keep your configuration, click **Done**.
- If you are editing an existing configuration and want to keep your changes, click **Save**.

Review the SP STS configuration

When you are finished with the **WS-Trust STS** task, you can review the configuration on its **Summary** screen.

- If you need to make any changes, click the heading over the information you want to edit.
- If you are creating a new connection and want to keep your configuration, click **Done**.
- If you are editing an existing configuration and want to keep your changes, click **Save**.

IdP-to-SP bridging

The IdP-to-SP Bridging links on the Server Configuration menu provide access to advanced federation settings.

This chapter covers:

- [Adapter-to-adapter mappings](#) on page 420
- [Token translator mappings](#) on page 423
- [Federation hub and authentication policy contracts](#) on page 89

 **Note:** The information in this chapter is presented from the viewpoint of an administrative user with “Admin” permissions (see [Account management](#) on page 108).

Adapter-to-adapter mappings

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 via an SP adapter on the same PingFederate server. Generally, these cases involve SaaS providers who may not support standards-based SSO but do provide proprietary SSO with “delegated authentication” (for example, Salesforce and Workday).

The mapping may also be used to enable the Google Apps Password Manager (available separately with the Google Apps Connector—see [Outbound provisioning for IdPs](#) on page 81).

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 (a loopback 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. Note that you may reuse instances that are also in use for connection configurations. For more information, see:

- [Manage IdP adapters](#) on page 275
- [Manage SP adapters](#) on page 331

Manage mappings

On the **Adapter-toAdapter Mappings** screen you can add, modify, or delete adapter-to-adapter mappings.

To add a mapping:

- Select the adapter Source Instance (IdP) and Target Instance (SP) and 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.

To edit a mapping:


- Click the mapping name.

To delete a mapping:

- Click **Delete** under Actions for the mapping and then click **Save**.

Assign a license group

Adapter-to-adapter mapping is considered a connection for licensing purposes. If your PingFederate license manages connections by groups, select a license group for this mapping configuration.

 **Note:** This screen is not displayed for unrestricted or other types of licenses.

To assign a License Group:

- Select the License Group from the drop-down list and click **Next**.

Configure attribute lookup for adapter-to-adapter mapping

Attribute sources are specific data store or directory locations containing information that may be required to fulfill the SP adapter contract.

This portion of the adapter-to-adapter configuration allows you to configure one or more data stores to look up attributes and to set up search parameters.

- If data-store lookup is *not* required, click **Next**.
- If data-store lookup is required, click **Add Attribute Source** and complete the setup steps (see [Choose a data store](#) on page 508).

To modify an attribute source configuration:

1. Click the attribute source Description link.
2. Click **Save** on the screen you change.



Note: Depending on what you change, you may need to modify dependent data in subsequent steps, as indicated. Click **Save** or **Done** when either of those options appears.

Identify target application (optional)

Use this screen to enter the name of the target application and the URL of the application icon, accessible through the IdP Adapter interface (`IdpAuthenticationAdapterV2`) in the PingFederate Java SDK. (For more information about the SDK, see [SDK developer's guide](#).) Both fields are optional.

To complete the configuration:

- Enter an Application Name and an Application Icon URL.

Define adapter-to-adapter contract fulfillment

The next step in this configuration is to map values from the IdP adapter into the attributes required by the SP adapter (the Adapter Contract).

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 screen](#) on page 501). (If the Expression selection is not listed, then the feature is not enabled—see [Enable and disable expressions](#) on page 497. For syntax and examples, see sections under [Construct OGNL expressions](#) on page 498.)

- LDAP/JDBC/Custom (if a data store is configured)

Values are returned from a data source. When you make this selection, the Value list is populated by the LDAP, JDBC, or Custom attributes you identified.



Note: PingFederate appends a description in parentheses for configured data store lookups (see [Configure attribute lookup for adapter-to-adapter mapping](#) on page 420).

- Expression (when enabled)

This option provides more complex mapping capabilities—for example, transforming incoming values into different formats (see [Attribute mapping expressions](#) on page 497). All of the variables available for text entries (see below) are also available for expressions.

- 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.

You can also enter values from your data store, when applicable, using this syntax:

```
${ds.attr-source-id.attribute}
```

where `attr-source-id` is the **Attribute Source Id** value (see [Choose a data store](#) on page 508) and `attribute` is any of the data store attributes you select.

To map attributes:

1. Choose a Source for each SP Adapter Contract attribute.

See the list under [Map each attribute to fulfill the Adapter Contract from one of these Sources](#) above.

2. Choose (or enter) a Value for each attribute.

All values must be mapped.

3. Click **Next** (or **Done**, if you are finished with the configuration, and then **Save** on the Manage Mappings screen).

Configure a default target URL (optional)

Use this screen to assign a default target URL for this adapter-to-adapter mapping. Entering a URL in the Default Target URL field overrides any SP Default URL SSO setting (see [Configure default URLs](#) on page 336).



Note: If specified, the adapter-to-adapter endpoint parameter `TargetResource` overrides any default target URL for an adapter-to-adapter mapping (see [System-services endpoints](#) on page 460).

Configure issuance criteria (optional)

Use this screen to define criteria PingFederate can evaluate to determine whether to issue an SP adapter token for a user (see [About token authorization](#) on page 74). This token authorization can be used to restrict who can SSO to protected resources.

To configure Issuance Criteria:

1. Select a source containing attributes for token authorization.

Associated attributes appear in the Attribute Name drop-down list:

- Adapter – Select to access attributes from the IdP Adapter.
- Context – Select to use values 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. If you want to use an expression for authentication, skip to Step 8 below.

- LDAP/JDBC/Custom (if a data store is configured) – Select to access attributes returned from a data source.



Note: PingFederate appends a description in parentheses for configured data store lookups (see [Configure attribute lookup for adapter-to-adapter mapping](#) on page 420)

- Mapped Attributes – Select to access the required attributes.

2. Select an attribute name.

3. Select the Condition you want to apply.



Note: Use one of the *multi-value* conditions when you want PingFederate to verify that an attribute contains or does not contain the specified value in its attribute list. Using a single-value condition when an attribute has multiple values causes the condition to evaluate consistently to false.

4. Enter the exact value in the **Value** field.

Wildcards (for partial matching) are not allowed.




Tip: You can use OGNL expressions to perform more complex evaluations, including partial matching.

5. Optional: Use the Error Result field to customize an error code or error message for use if the token authorization fails.

Error results are handled in one of two ways:


- **Redirect** – 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 (for more information on this and other user-facing templates, see [Customizable user-facing screens](#) on page 126).

-  **Note:** Using an error code in the Error Result field allows the error template or a web application to process the error result in a variety of ways—for example, a redirect to another page or e-mail to an administrator.


If you leave this field blank, a default ACCESS_DENIED error result is used if the authorization fails.

6. Click **Add**.

7. Repeat the steps above to add additional attributes, as needed, for each authorization criteria.

-  **Note:** All criteria must pass in order for a user to be authorized.


8. Optional: Click **Show Advanced Criteria** to configure more complex evaluations using OGNL expressions.

-  **Note:** Expressions must be enabled for the **Show Advanced Criteria** button to appear (see [Enable and disable expressions](#) on page 497).


- Use the in-line editor box to enter the OGNL expression.

For more information about OGNL, see [Attribute mapping expressions](#) on page 497.

- Use the Error Result box to enter an error code or an error message for use if authorization fails (see step 5 for more information).

-  **Note:** If the OGNL expression resolves to a string value (rather than true or false), then the returned value overrides the Error Result.

- Click **Add**. Continue to add expressions, as needed.
- Click the **Test** link to input values and test the resulting output for the expression.

-  **Note:** For more information on testing OGNL expressions, see [Using the OGNL edit screen](#) on page 501. For syntax and examples, see sections under [Construct OGNL expressions](#) on page 498.

9. Click **Next** (or **Done**, if you are finished with the configuration, and then **Save** on the Manage Mappings screen).

To edit Issuance Criteria:

- Click **Edit** under Actions for the criteria.

To delete Issuance Criteria:

- Click **Delete** under Actions for the criteria and then click **Save**.

Review the summary screen

When you have finished configuring Adapter-to-Adapter Mapping, you can review the configuration on the Summary screen.

If you need to make any changes, click the heading over the information you want to edit. When you are finished, click **Done** and then **Save** on the Manage Mappings screen.

Token translator mappings

This configuration is provided for use cases in which the PingFederate WS-Trust STS exchanges one type of security token for another without requiring a SAML token to be generated in between (see [About WS-Trust STS](#) on page 58). Use this configuration, for example, to convert a user's Kerberos token to a third-party proprietary 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 [Enable the WS-Trust protocol](#) on page 400). Also, be sure to configure the required token-translator instances. Note that you may reuse instances that are also in use for STS connection configurations. For more information, see:

- [Manage token processors](#) on page 401
- [Manage token generators](#) on page 412


Manage token mappings

On the **Token Translator Mappings** screen you can add, modify, or delete token-to-token mappings.

To reach this screen:

1. Click **Server Configuration** on the Main Menu.
2. Click **Token Translator Mappings** under IdP-to-SP Bridging.

To add a mapping:

- Select the token-processor Source Instance and the token-generator Target Instance and 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 multiple IdP token processors and/or SP token generators are configured, you must provide the `TokenProcessorId` and/or the `TokenGeneratorId` query parameter(s) in the request (see [System-services endpoints](#) on page 460).

To edit a mapping:

- Click the mapping name.

To delete a mapping:

- Click **Delete** under Actions for the mapping and then click **Save**.

Configure attribute lookup for token mapping

Attribute sources are specific data store or directory locations containing information that may be required to fulfill a token-generator contract.

This optional portion of the configuration allows you to configure one or more data stores to look up attributes and to set up search parameters.

- If data-store lookup is *not* required, click **Next**.
- If data-store lookup is required, click **Add Attribute Source** and complete the setup steps (see [Choose a data store](#) on page 508).

To modify an attribute source configuration:

1. Click the attribute source Description link.
2. Click **Save** on the screen you change.



Note: Depending on what you change, you may need to modify dependent data in subsequent steps, as indicated. Click **Save** or **Done** when either of those options appears.

Define token contract fulfillment

The next step in this configuration is to map values from the token processor into the attributes required by the token generator (the Token Generator Contract).


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.

- 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 screen](#) on page 501). (If the Expression selection is not listed, then the feature is not enabled—see [Enable and disable expressions](#) on page 497. For syntax and examples, see sections under [Construct OGNL expressions](#) on page 498.)

- LDAP/JDBC/Custom (if a data store is configured)

Values are returned from a data source. When you make this selection, the Value list is populated by the LDAP, JDBC, or Custom attributes you identified.

 **Note:** PingFederate appends a description in parentheses for configured data store lookups (see [Fulfillment by data store queries](#) on page 506).

- Expression (when enabled)

This option provides more complex mapping capabilities—for example, transforming incoming values into different formats (see [Attribute mapping expressions](#) on page 497). All of the variables available for text entries (see below) are also available for expressions.

- 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 data store, when applicable, using this syntax:

```
${ds.attr-source-id.attribute}
```

where `attr-source-id` is the **Attribute Source Id** value (see [Fulfillment by data store queries](#) on page 506) and `attribute` is any of the data store attributes you select.

To map attributes:

1. Choose a Source for each Token Generator Contract attribute.

See the list under [Map each attribute to fulfill the Token Generator Contract from one of these Sources](#): above.

2. Choose (or enter) a Value for each attribute.

All values must be mapped.

3. Click **Next** (or **Done**, if you are finished with the configuration, and then **Save** on the Manage Mappings screen).

Specify token issuance criteria


Use this optional configuration to define criteria PingFederate can evaluate to determine whether to issue a security token (see [About token authorization](#) on page 74). This token authorization can be used to restrict who can access protected Web Services.

To configure Issuance Criteria:

1. Select a source containing attributes for token authorization.

Associated attributes appear in the Attribute Name drop-down list:


- Context – Select to use values returned from the context of the transaction at runtime.

 **Note:** The HTTP Request and STS SSL Client Certificate Chain selections are retrieved as Java objects rather than text. For this reason, OGNL expressions are more appropriate to evaluate and return values. If you want to use an expression for authentication, skip to Step 8 below.

Choose Expression and then click **Edit** to enter an expression (see [Using the OGNL edit screen](#) on page 501). (If the Expression selection is not listed, then the feature is not enabled—see [Enable and disable expressions](#) on page 497.

For syntax and examples, see sections under [Construct OGNL expressions](#) on page 498.)


- LDAP/JDBC/Custom (if a data store is configured) – Select to access attributes returned from a data source.


 **Note:** PingFederate appends a description in parentheses for configured data store lookups (see [Fulfillment by data store queries](#) on page 506).

- Mapped Attributes – Select to access the required attributes from the token contract.
- Token – Select to access attributes from the IdP token processor.

2. Select an attribute name.


3. Select the Condition you want to apply.

 **Note:** Use one of the *multi-value* conditions when you want PingFederate to verify that an attribute contains or does not contain the specified value in its attribute list. Using a single-value condition when an attribute has multiple values causes the condition to evaluate consistently to false.

 **Important:** When using the STS SSL Client Certificate's Subject DN attribute, you must select one of the following conditions: equal to DN, not equal to DN, multi-value contains DN, or multi-value does not contain DN. These operators take into account minor differences that might be present in the DN syntax (for example, whether spaces are present after commas).

4. Enter the exact value in the **Value** field.


Wildcards (for partial matching) are not allowed.

 **Tip:** You can use OGNL expressions to perform more complex evaluations, including partial matching.

5. Optional: Use the Error Result field to customize an error code or error message for use if the token authorization fails.

Errors result in a SOAP message returned. 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](http://www.w3.org/TR/2000/NOTE-SOAP-20000508/#_Toc478383507) (www.w3.org/TR/2000/NOTE-SOAP-20000508/#_Toc478383507).

 **Note:** 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.


If you leave this field blank, the user sees a default ACCESS_DENIED error result at runtime if the authorization fails.

6. Click **Add**.

7. Repeat the steps above to add additional attributes, as needed, for each authorization criteria.

 **Note:** All criteria must pass in order for a user to be authorized.


8. Optional: Click **Show Advanced Criteria** to configure more complex evaluations using OGNL expressions.

 **Note:** Expressions must be enabled for the Show Advanced Criteria button to appear (see [Enable and disable expressions](#) on page 497).


- Use the in-line editor box to enter the OGNL expression.

For more information about OGNL, see [Attribute mapping expressions](#) on page 497.

- Use the Error Result box to enter an error code or an error message for use if authorization fails (see step 5 above for more information).

 **Note:** If the OGNL expression resolves to a string value (rather than true or false), then the returned value overrides the Error Result.

- Click **Add**. Continue to add expressions, as needed.
- Click the **Test** link to input values and test the resulting output for the expression.

 **Note:** For more information on testing OGNL expressions, see [Using the OGNL edit screen](#) on page 501.

9. Click **Next** (or **Done**, if you are finished with the configuration, and then **Save** on the Manage Mappings screen).

To edit Issuance Criteria:

- Click **Edit** under Actions for the criteria.

To delete Issuance Criteria:

- Click **Delete** under Actions for the criteria and then click **Save**.

Review the token mapping summary screen

When you have finished configuring token-to-token mapping, you can review the configuration on the Summary screen.

If you need to make any changes, click the heading over the information you want to edit. When you are finished, click **Done** and then **Save** on the Manage Mappings screen.

Bundled adapters

PingFederate comes bundled with a set of adapters:

- [Kerberos Adapter](#) on page 427
- [HTML Form Adapter](#) on page 430
- [HTTP Basic Adapter](#) on page 438
- [OpenToken Adapter](#) on page 440
- [Composite Adapter](#) on page 445

Kerberos Adapter

The integrated Kerberos Adapter provides a seamless SSO experience for Windows clients by authenticating SSO requests using the Kerberos v5 protocol against Active Directory (AD) domains.

When the PingFederate® IdP server receives an authentication request for SP-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.



Note: 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 one or more instances of it alongside with the Kerberos Adapter.

Authentication mechanism assurance

For the purpose of protecting resources based on login 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 logs on using a certificate-based login method, such as a smart-card login. For example, you can restrict access to sensitive resources to users whom log 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 workflow to verify that the SIDs attribute contains the SID value associated with the required login method. If the SIDs attribute does not contain the specified SID value, the request is denied.



Note: 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](https://technet.microsoft.com/en-us/library/dd378897%28v=ws.10%29.aspx) from Microsoft (technet.microsoft.com/en-us/library/dd378897%28v=ws.10%29.aspx).

Configure the Kerberos Adapter


1. If you have not already done so, log on to the PingFederate® administrative console.
2. Go to the **IdP Configuration > Adapters** screen.
3. On the **Manage IdP Adapter Instances** screen, click **Create New Instance**.
4. On the **Type** screen, enter an **Instance Name** and **Instance Id**, select **Kerberos Adapter** from the list, and then click **Next**.

The **Instance Id** may not contain spaces or underscores.


5. Optional: Select a **Parent Instance** from the 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 overridden. You can specify overrides during the rest of the setup.

6. Click **Next**.

 **Note:** If this is a child instance, select the override check box related to the settings you want to modify.

7. On the **IdP Adapter** screen, select the **Domain/Realm Name** for your Windows domain. If the domain or realm you want does not appear, click **Manage Active Directory Domains/Kerberos Realms** to add it (see [Configure Active Directory domains or Kerberos realms](#) on page 203).
8. Optional: 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 screen 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 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.


9. Optional: Click **Show Advanced Fields** and make any desired changes for the following settings.

Field	Descriptions
Error Template (optional)	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</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 should authentication fail.
Authentication Context Value (optional)	This may 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 (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: <ul style="list-style-type: none"> • <code>urn:oasis:names:tc:SAML:1.0:am:unspecified</code> for SAML 1.x

Field	Descriptions
	<ul style="list-style-type: none"> urn:oasis:names:tc:SAML:2.0:ac:classes:unspecified for SAML 2.0 <p>As needed, the authentication context can be overridden by either an instance of the Requested AuthN Context Authentication Selector or the <code>SAML_AUTHN_CTX</code> attribute in the SAML attribute contract. (The latter takes precedence.)</p>

10. Click **Next**.

11. On the **Adapter Attributes** screen, select **Username** (and optionally **Domain/Realm Name**) to be used in constructing a unique identifier (Pseudonym) for account linking at the SP.

 **Note:** A selection is required regardless of whether 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 (for example, email) to prevent the same pseudonym from being assigned to multiple users.


You can also choose to mask the values of any or all attributes that PingFederate logs from the adapter at runtime.

If OGNL expressions might be used to map derived values into outgoing assertions and you want those values masked, select the related check box under the Attribute list.

12. Click **Next**.

13. On the **Summary** screen, click **Done**.


14. On the **Manage IdP Adapter Instances** screen, click **Save**.


 **Important:** You must click **Save** if you want to retain the adapter configuration.

Configure end-user browsers

Follow the steps in the following sections to configure client-side browsers at your site in order to use the Kerberos Adapter to authenticate users:

- [Configure Microsoft Internet Explorer](#) on page 429
- [Configure Mozilla Firefox](#) on page 430


 **Note:** The client-side configuration requires the **Base URL** value of your PingFederate® environment, which can be found in **Server Configuration > Server Settings > Federation Info**.

 **Important:** If the browsers are not properly configured, users may be prompted to authenticate manually using their network credentials or fail to SSO to the service providers.

Configure Microsoft Internet Explorer


To configure Internet Explorer (9 or higher) for Kerberos authentication, review the following settings in Internet Options.

1. Add the base URL to **Local intranet**.

 **Note:** This step may be skipped if the base URL (<pf-idp.domain.name>) is internal and not fully qualified. For example, if it is `pingfederate`, you can skip this step. However, if <pf-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.

- Close all Internet Explorer tabs and windows.
- Open **Control Panel > Internet Options**.
- Click the **Security** tab.
- Select **Local intranet** and click **Sites**.
- Click **Advanced**.
- Enter the base URL (for example, `www.example.com`), and then click **Add**.
- Click **Close**, and then click **OK** to return to the Security tab.

2. Verify **Automatic logon only in the Intranet zone** is selected.
 - a) Under the Security tab, select **Local intranet** and 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.

 - a) Click the **Connections** tab.
 - b) Click **LAN settings**.
 - c) Verify the **Use a proxy server for your LAN ...** check box is selected, and then click **Advanced**.
 - d) Enter the base URL in the **Exceptions** field, and then click **OK**.
 - e) Click **OK** to return to the Connections tab.
4. Verify **Enable Integrated Windows Authentication** is selected.
 - a) Click the **Advanced** tab.
 - b) Verify **Enable Integrated Windows Authentication** is selected in the **Settings** pane.
5. Click **OK** to close Internet Options.

Configure Mozilla Firefox

To configure Firefox for Kerberos authentication, configure Firefox as follows:

1. Start Firefox.
2. Open a new tab, and then enter `about:config` in the address bar.
3. Search for the `network.negotiate-auth.trusted-uris` preference name.
4. Double-click to modify its value to include the base URL of your PingFederate® environment (for example, `www.example.com`).
5. Click **OK** and close the `about:config` tab.
6. Optional: Exit Firefox.

HTML Form Adapter

Initial user authentication is normally handled outside of the PingFederate® server using an application or IdM system logon module. PingFederate's adapter and application agents are typically used to integrate with these local authentication mechanisms.

PingFederate packages an HTML Form Adapter that delegates user authentication to a configured password credential validator. This authentication mechanism validates credentials based on either an LDAP directory or a simple username validator that authenticates credentials maintained by PingFederate.

On the IdP side, when the PingFederate IdP server receives an authentication request for SP-initiated SSO or the user clicks a link for IdP-initiated SSO, the IdP server invokes the HTML Form Adapter and, if not already authenticated, prompts the user for local IdP credentials. The credentials are then validated using the designated password credential validator and, if validated, a SAML assertion is generated.

The HTML Form Adapter allows you to customize a different login page for each configured adapter instance. You can define a logout path or a logout redirect page. You can also enable users to change or reset their network passwords or redirect users to a company-hosted password management system.

PingFederate 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 **Challenge Retries** is three. If you prefer a higher value, consider reviewing the account lockout policy of your directory servers first. For example, if the directory server's account lockout threshold is set to five and the **Challenge Retries** in the HTML Form Adapter is also set to five (or a higher value), the fifth single sign-on attempt could potentially lock the user accounts on the directory server.

The logout period is controlled by the Account Locking Service. For configuration information, see [Account lockout protection](#).

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.)

Configure the HTML Form Adapter

1. Click **IdP Configuration > Adapters** to open the **Manage IdP Adapter Instances** screen.
2. On the **Manage IdP Adapter Instances** screen, click **Create New Instance** to start the **Create Adapter Instance** workflow.
3. On the **Type** screen, configure the basics of this adapter instance.
 - a) Enter the required information and select the adapter type from the 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 screens.

4. On the **IdP Adapter** screen, 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) Select a Password Credential Validator instance from the list and click **Update**.





Add as many validators as necessary. Click **Move up** and **Move down** 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 is able to 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 lockout those accounts in their source locations.

- d) Enter values for the adapter configuration, as described below.

Property	Description
Challenge Retries (Required)	The account lockout threshold for this adapter instance. When the number of login failures reaches this threshold, the user is locked out for one minute. The default value is 3.
Session State	How the session state is maintained between adapters of the same type. When Globally is selected, sessions are shared among other HTML Form Adapter instances that use the same Session State field value (' Globally '). When Per Adapter is selected, a session is tied to a specific adapter instance. When None is selected, a session is not maintained. The default selection is Per Adapter .

Property	Description
Session Timeout	<p>The number of idle minutes before a session times out based on inactivity. If left blank, the lifetime falls back on the Session Max Timeout field value. Ignored if None is selected for the Session State field.</p> <p>The default value is 60 (minutes).</p>
Session Max Timeout	<p>The maximum lifetime (in minutes) before a session expires regardless if the Session Timeout field value has been reached, or not. Ignored if None is selected for the Session State field.</p> <p>The default value is 480 (minutes, which translates to 8 hours).</p> <p> Note: This setting sets a maximum lifetime, subject to inactivity timeout. Consider the following examples:</p> <p>A user initiated an SSO request at 9 a.m. and has not made another SSO request since then. At 10 a.m., the session times out based on inactivity (based on the default Session Timeout field value of 60 minutes).</p> <p>Another user initiated an SSO request at 9 a.m. and has been making SSO requests every hour at least once. This session does not time out because the user has been actively making SSO requests; however, the session does expire at 5 p.m. (based on the default Session Max Timeout default value of 8 hours).</p> <p>If you leave both the Session Max Timeout and Session Timeout fields blank, sessions do not expire (until PingFederate® restarts or the sessions are cleaned up by another means).</p> <p>If you leave the Session Max Timeout field blank but set a value for the Session Timeout field, sessions do not expire until they time out based on inactivity.</p> <p> Tip: Session information is stored in the PF cookie. By default, the PF cookie is a session cookie and is typically removed by the web browsers on exit.</p> <p>You can optionally extend the lifetime of the PF cookie by editing the <code>session-cookie-config.xml</code> file (located in the <code><pf_install>/pingfederate/server/default/data/config-store</code> directory).</p>
Allow Password Changes	<p>Enables or disables the ability for users to change their network password using this adapter instance. Select the check box to enable the change password functionality.</p> <p>This check box is not selected by default.</p> <p> Note: The LDAP Username Password Credential Validator (PCV) and the PingOne® directory PCV are currently the only PCVs bundled with PingFederate that support the change password feature.</p> <p>Additionally, if the underlying PCV instance is connected to a Microsoft Active Directory (AD) LDAP server, you <i>must</i> secure the data store connection to your AD LDAP server using LDAPS; AD requires this level of security to allow password changes.</p> <p> Tip: When selected, if an IdP-initiated SSO request or an adapter-to-adapter request includes the <code>ChangePassword</code> parameter with a value of <code>true</code>, the user is redirected to the Change Password form immediately.</p>

Property	Description
	<p>For more information, see IdP endpoints on page 448 or System-services endpoints on page 460.</p>
Password Management System	<p>The URL for redirecting users to a company-specific password management system to change their password.</p> <p>This field is left blank by default.</p>
Enable 'Remember My Username'	<p>Allows users to store their username as a cookie when authenticating with this adapter. Once stored, the username in the login form is pre-populated for subsequent transactions. Select the check box to enable the cookie functionality.</p> <p>This check box is not selected by default.</p> <p> 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.</p>
Change Password Email Notification	<p>When selected, an email notification is sent to 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 PCV.</p> <p>This check box is not selected by default.</p> <p> Note: This option requires the selection of the Allow Password Changes check box and a pre-configured email server on the Server Configuration > Server Settings > Runtime Notification screen.</p> <p>In addition, the LDAP PCV is the only PCV bundled with PingFederate that supports this email notification feature.</p>
Show Password Expiring Warning	<p>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.</p> <p>This check box is not selected by default.</p> <p> Note: This option requires the selection of the Allow Password Changes check box. (Both check boxes are not selected by default.)</p> <p>In addition, the LDAP PCV is currently the only PCV bundled with PingFederate that supports the password expiring warning feature.</p>
Password Reset Type	<p>Select one of the following methods for self-service password reset (SSPR).</p> <ul style="list-style-type: none"> • Email One-Time Link: users receive an email with a URL to reset their password. • Email One-Time Password: users receive an email with a one-time password (OTP) to reset their password. <p>These options require a pre-configured email server on the Server Configuration > Server Settings > Runtime Notification screen.</p> <ul style="list-style-type: none"> • PingID: users are prompted to follow the PingID® authentication flow to reset their password.

Property	Description
	<p>Ensure the PingID Username Attribute field in the selected LDAP PCV instance is configured; otherwise, users will not be able to reset their password.</p> <ul style="list-style-type: none"> • Text Message: users receive a text message notification with an OTP to reset their password. <p>Ensure the SMS Attribute field in the selected LDAP PCV instance is configured; otherwise, users will not receive text message notification for password reset.</p> <p>This option also requires SMS provider settings (see <i>the last sub step</i>).</p> <ul style="list-style-type: none"> • None: users cannot reset password through this HTML Form Adapter instance. This is the default selection. <p> Tip: Provided the email server configuration on the Server Configuration > Server Settings > Runtime Notification > Email Server Settings screen is valid, an email notification is sent to the user who has successfully reset the password through the HTML Form Adapter. The destination is the user's email address, specifically the value of the attribute defined by the Mail Attribute setting in the LDAP PCV.</p> <p> Note: To enable password reset, you must also select the Allow Password Changes check box.</p> <p>In addition, the LDAP PCV is the only PCV bundled with PingFederate that supports self-service password reset.</p>
Account Unlock	<p>Select this check box to enable self-service account unlock using this adapter instance.</p> <p>This check box is not selected by default.</p> <p> Note: You must also select a Password Reset Type type other than None (and therefore the Allow Password Changes check box as well) because the initiating user must prove ownership of the account through the password reset flow.</p> <p>Unlike password reset, 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 Active Directory. If the underlying data store is connected to an Oracle Directory Server, users can only unlock their account by changing their current password through the password reset flow.</p>
	<p>e) Click Show Advanced Fields to review or modify default values.</p> <p>For information about each field, see <i>HTML Form Adapter advanced fields</i> on page 435.</p> <p>f) If you have chosen Text Message as the password reset type, click Manage SMS Provider Settings to configure the SMS provider through which PingFederate can send text message notifications to the users.</p> <p>5. On the Extended Contract screen, configure additional attributes for this adapter instance as needed.</p> <p>6. On the Adapter Attributes screen, configure the pseudonym and masking options.</p> <p> Note: The Override Attributes check box in this screen reflects the status of the override option in the Extended Contract screen.</p> <p>a) Select the check box under Pseudonym for the user identifier of the adapter and optionally for the other attributes, if available.</p> <p>This selection is used if any of your SP partners use pseudonyms for account linking.</p>





Note: A selection is required regardless of whether 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 (for example, email) to prevent the same pseudonym from being assigned 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** screen, configure the adapter contract for this instance with the following optional workflows:
 - Configure one or more data sources for data store queries.
 - Fulfill adapter contract with values from the adapter (the default), data store queries (if configured), context of the request, text, or expressions (if enabled).
 - Set up the token authorization workflow to validate one or more criteria prior to the issuance of the adapter contract.
 8. On the **Summary** screen, review your configuration, modify as needed, and click **Done** to exit the **Create Adapter Instance** workflow.
 9. On the **Manage IdP Adapter Instances** screen, click **Save** to retain the configuration of the adapter instance.
If you want to exit without saving the configuration, click **Cancel**.

HTML Form Adapter advanced fields

Property	Description
Login Template (Required)	<p>The HTML template to prompt the users for their credentials. PingFederate allows each configured adapter instance to use a different login page template.</p> <p>The default template file is <code>html.form.login.template.html</code>.</p> <p>(Unless otherwise stated, all template files are located in the <code><pf_install>/pingfederate/server/default/conf/template</code> directory.)</p>
Logout Path	<p>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 URL is <code><PingFederate base URL>/ext/<Logout Path></code>.</p> <p>Available primarily for use cases where the partner 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.</p> <p>This field is left blank by default.</p> <p> Note: If specified, the path must be unique across all HTML Form Adapter instances, including child instances.</p>
Logout Redirect	<p>The landing page at the SP after successful IdP logout (applicable only when the Logout Path field is configured).</p> <p>This field is left blank by default.</p>
Logout Template	<p>The HTML template to be 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.</p> <p>The default template file is <code>idp.logout.success.page.template.html</code>.</p>
Change Password Template	<p>The HTML template to prompt the users to change their password. PingFederate allows each configured adapter instance to use a different change password template.</p>

Property	Description
	The default template file is <code>html.form.change.password.template.html</code> .
Change Password Message Template	<p>The HTML template to be displayed when a user has successfully changed the password through the HTML Form Adapter.</p> <p>The default template file is <code>html.form.message.template.html</code>.</p>
Password Management System Message Template	<p>The HTML template to notify the users that they are being redirected to a password management system to change their password.</p> <p>The default template file is also <code>html.form.message.template.html</code>.</p>
Change Password Email Template	<p>The HTML message template to notify the user via email that the password has been changed or reset successfully through the HTML Form Adapter.</p> <p>The default template file is <code>message-template-end-user-password-change.html</code>, which is located in the <code><pf_install>/pingfederate/server/default/conf/template/mail-notifications</code> directory.</p>
Expiring Password Warning Template	<p>The HTML template to warn the users about approaching the password expiry day.</p> <p>The default template file is <code>html.form.password.expiring.notification.template.html</code>.</p>
Threshold for Expiring Password Warning	<p>The threshold (in days) to start warning the user about approaching the password expiry day.</p> <p>The default value is 7 (days).</p>
Snooze Interval for Expiring Password Warning	<p>The amount of time (in hours) to delay the next warning after the user has chosen to change the password later.</p> <p>The default value is 24 (hours).</p>
Login Challenge Template	<p>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 is an example of where it could be used.</p> <p>The default template file is <code>html.form.login.challenge.template.html</code>.</p>
'Remember My Username' Lifetime	<p>Number of days the cookie remains valid. Enter the number of days you want the username remembered in a cookie.</p> <p>The cookie lifetime is reset upon each successful login in which the Enable 'Remember My Username' check box is selected.</p> <p>The default value is 30 (days).</p> <p> 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.</p>
Allow Username Edits During Chaining	<p>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 <code>login_hint</code> parameter, the username in the login form is pre-populated; users are not allowed to edit their usernames.</p> <p>Select this check box if you want to allow users to edit the pre-populated username in the login form.</p> <p>This check box is not selected by default.</p>

Property	Description
	 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.
Track Authentication Time	<p>When selected, the time of authentication for each user is tracked and could be utilized by applicable use cases. For example, if an OAuth client sends an authorization request with a <code>max_age</code> 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 <code>max_age</code> value.</p> <p>This check box is selected by default.</p>
Post-Password Change Re-Authentication Delay	<p>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 should wait prior to the reauthentication attempt.</p> <p>The default value is 0, which is the minimum value. The maximum value is 60000 (one minute).</p>

Advanced fields for self-service password reset

Property	Description
Password Reset Username Template	<p>The HTML template to prompt the user to enter a username for password reset.</p> <p>This template is applicable for all password reset types (other than None).</p> <p>The default template file is <code>forgot-password.html</code>.</p>
Password Reset Code Template	<p>The HTML template to prompt the user to enter the one-time password (OTP) or the code for password reset.</p> <p>This template is applicable when the password reset type is Email One-Time Password or Text Message.</p> <p>The default template file is <code>forgot-password-resume.html</code>.</p>
Password Reset Template	<p>The HTML template to prompt the user to define a new password.</p> <p>This template is applicable for all password reset types (other than None).</p> <p>The default template file is <code>forgot-password-change.html</code>.</p>
Password Reset Error Template	<p>The HTML template to notify the user that the password reset attempt has failed.</p> <p>This template is applicable for all password reset types (other than None).</p> <p>The default template file is <code>forgot-password-error.html</code>.</p>
Password Reset Success Template	<p>The HTML template to notify the user that the password reset attempt has succeeded.</p> <p>This template is applicable for all password reset types (other than None).</p> <p>The default template file is <code>forgot-password-success.html</code>.</p>
Account Unlock Template	<p>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.</p> <p>The default template file is <code>account-unlock.html</code>.</p>
OTP Length	<p>The number of characters in the one-time password for password reset.</p>


Property	Description
	The default value is 8.
OTP Time to Live	The validity (in minutes) for the one-time password or the one-time link. The default value is 10 (minutes).
PingID Properties	For self-service password reset using PingID®, follow these steps to upload the PingID settings file to the HTML Form Adapter instance: <ol style="list-style-type: none"> 1. Sign on to the PingOne® admin portal. 2. Go to the Setup > PingID > Client Integration screen. 3. Download the settings file (<code>pingid.properties</code>). 4. If the Integrate with PingFederate and Other Clients check box is not selected, select it and then click Save. 5. Close the PingOne admin portal. 6. Come back to the PingFederate administrative console and upload the <code>pingid.properties</code> file to the PingID Properties advanced field on the IdP Adapter screen.

Manage 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.

1. Go to the **IdP Configuration > Adapters** screen.
2. Select any HTML Form Adapter instance, in which the selected password reset type is **Text Message**.
3. On the **SMS Provider Settings** screen, 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.

 **Tip:** For additional information about each field or Twilio, please refer to support.twilio.com.

Once saved, these SMS provider settings apply to all services using text message notifications.

HTTP Basic Adapter

Initial user authentication is normally handled outside of the PingFederate® server using an application or IdM system logon module. PingFederate's adapter and application agents are typically used to integrate with these local authentication mechanisms.

PingFederate packages an HTTP Basic Adapter that delegates user authentication to a configured password credential validator. This authentication mechanism validates credentials based on either an LDAP directory or a simple username validator that authenticates credentials maintained by PingFederate.

On the IdP side, when the PingFederate IdP server receives an authentication request for SP-initiated SSO or the user clicks a link for IdP-initiated SSO, the IdP server invokes the HTTP Basic Adapter and, if not already authenticated, prompts the user for local IdP credentials. The credentials are then validated using the designated password credential validator and, if validated, a SAML assertion is generated.

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.)

Configure the HTTP Basic Adapter

1. Click **IdP Configuration > Adapters** to open the **Manage IdP Adapter Instances** screen.
2. On the **Manage IdP Adapter Instances** screen, click **Create New Instance** to start the **Create Adapter Instance** workflow.
3. On the **Type** screen, configure the basics of this adapter instance.
 - a) Enter the required information and select the adapter type from the 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 screens.

4. On the **IdP Adapter** screen, 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) Select a Password Credential Validator instance from the list and click **Update**.

Add as many validators as necessary. Click **Move up** and **Move down** 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 is able to 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 lockout those accounts in their source locations.

- d) Enter values for the adapter configuration, as described below.

Property	Description
Realm (Required)	The name of a protected area. The value of this field is sent as a part of 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 re-authenticate.
Challenge Retries (Required)	The number of attempts allowed for password authentication. The default value is 3.

5. On the **Extended Contract** screen, configure additional attributes for this adapter instance as needed.
6. On the **Adapter Attributes** screen, configure the pseudonym and masking options.



Note: The **Override Attributes** check box in this screen reflects the status of the override option in the **Extended Contract** screen.

- 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 SP partners use pseudonyms for account linking.



Note: A selection is required regardless of whether 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 (for example, email) to prevent the same pseudonym from being assigned 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** screen, configure the adapter contract for this instance with the following optional workflows:
 - Configure one or more data sources for data store queries.
 - Fulfill adapter contract with values from the adapter (the default), data store queries (if configured), context of the request, text, or expressions (if enabled).
 - Set up the token authorization workflow to validate one or more criteria prior to the issuance of the adapter contract.
 8. On the **Summary** screen, review your configuration, modify as needed, and click **Done** to exit the **Create Adapter Instance** workflow.
 9. On the **Manage IdP Adapter Instances** screen, click **Save** to retain the configuration of the adapter instance.

If you want to exit without saving the configuration, click **Cancel**.

OpenToken Adapter

In order to transfer identity and other user information between the PingFederate® server and an end application, the product 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 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 has the option to provide an authentication context to the SP by including the `authnContext` attribute with the desired value in the secure token. Standard URIs are defined in the SAML specifications (see the OASIS documents [oasis-sslc-saml-core-1.1.pdf](#) and [saml-authn-context-2.0-os.pdf](#)).

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 this topic to set up the OpenToken Adapter to use with your applications.

The following figure shows a basic IdP-initiated SSO scenario using PingFederate with the Java Integration Kit on both sides of an identity federation.

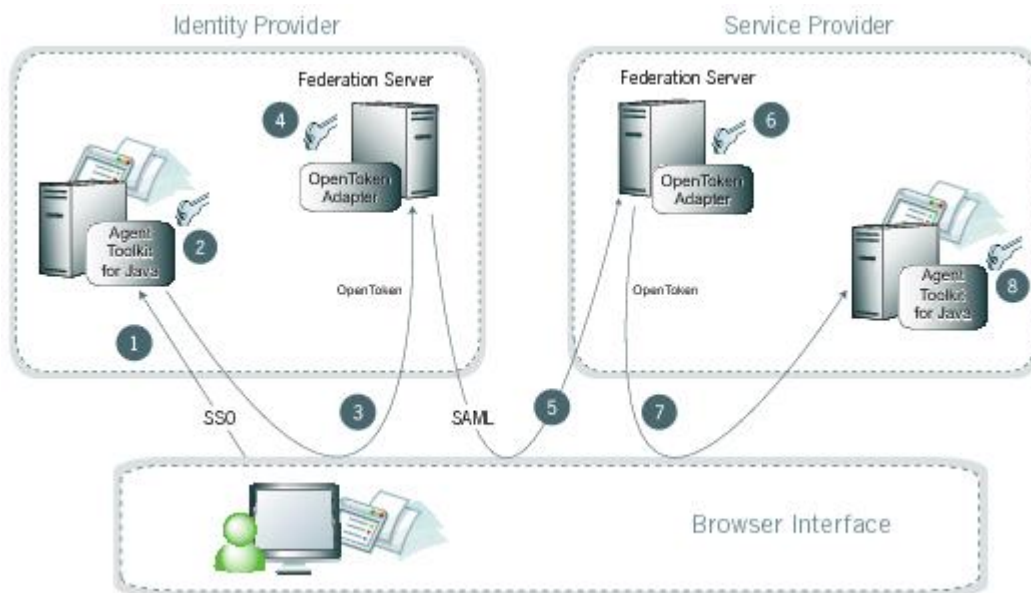


Figure 21: SP-Initiated SSO: POST/POST

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 Security Assertion Markup Language (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.


Configure the OpenToken IdP Adapter

1. Click **IdP Configuration > Adapters** to open the **Manage IdP Adapter Instances** screen.
2. On the **Manage IdP Adapter Instances** screen, click **Create New Instance** to start the **Create Adapter Instance** workflow.
3. On the **Type** screen, configure the basics of this adapter instance.
 - a) Enter the required information and select the adapter type from the 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 screens.

4. On the **IdP Adapter** screen, configure your OpenToken IdP Adapter instance. Refer to the on-screen field descriptions and the following table for more information.

 **Note:** These values are dependent on your developer's implementation.

Field	Description
Password	The password to use for generating the encryption key. It is also referred as the <i>shared secret</i> .
Confirm Password (Required)	
Authentication Service (Required)	The URL to which the user is redirected for an SSO event. This URL is part of an external application, which performs user authentication.
Click Show Advanced Fields to review the following settings.	
Transport Mode	How the token is transported to and from the application, either via a query parameter, a cookie (default), or as a form POST.
Token Name (Required)	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 (<code>opentoken</code>) 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, preceded by a period (for example, <code>.example.com</code>). 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).
Session Lifetime (Required)	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).
Not Before Tolerance (Required)	The amount of time (in seconds) to allow for clock skew between servers. Valid range is 0 to 3600. The default value is 0.
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.

Field	Description
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. The check box is not selected by default.  Note: Selecting this option may 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. We recommend to update your applications with the latest version of the agent. We recommend not to change the value of this flag.


5. On the **Actions** screen, click **Download** under **Action Invocation Link**, and then 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 IdP application. Refer to the documentation of your respective integration kit for more information.

6. On the **Extended Contract** screen, configure additional attributes for this adapter instance as needed.


Note that the OpenToken IdP Adapter always extends the core contract with an attribute `userId` and fulfills it with the value of `subject` for backward compatibility reason.

7. On the **Adapter Attributes** screen, configure the pseudonym and masking options.

 **Note:** The **Override Attributes** check box in this screen reflects the status of the override option in the **Extended Contract** screen.

- 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 SP partners use pseudonyms for account linking.

 **Note:** A selection is required regardless of whether 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 (for example, email) to prevent the same pseudonym from being assigned 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** screen, configure the adapter contract for this instance with the following optional workflows:

- Configure one or more data sources for data store queries.
- Fulfill adapter contract with values from the adapter (the default), data store queries (if configured), context of the request, text, or expressions (if enabled).
- Set up the token authorization workflow to validate one or more criteria prior to the issuance of the adapter contract.

9. On the **Summary** screen, review your configuration, modify as needed, and click **Done** to exit the **Create Adapter Instance** workflow.

10. On the **Manage IdP Adapter Instances** screen, click **Save** to retain the configuration of the adapter instance.

If you want to exit without saving the configuration, click **Cancel**.

Configure the OpenToken SP Adapter

1. Click **SP Configuration > Adapters** to open the **Manage SP Adapter Instances** screen.
2. On the **Manage SP Adapter Instances** screen, click **Create New Instance** to start the **Create Adapter Instance** workflow.
3. On the **Type** screen, configure the basics of this adapter instance.
 - a) Enter the required information and select the adapter type from the 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 screens.

4. On the **Instance Configuration** screen, configure your OpenToken SP Adapter instance. Refer to the on-screen field descriptions and the following table for more information.



Note: These values are dependent on your developer's implementation.

Field	Description
Password	The password to use for generating the encryption key. It is also referred as the <i>shared secret</i> .
Confirm Password (Required)	
Click Show Advanced Fields to review the following settings.	
Transport Mode	How the token is transported to and from the application, either via 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 (<code>opentoken</code>) 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 an 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, preceded by a period (for example, <code>.example.com</code>). 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).
Session Lifetime (Required)	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).

Field	Description
Not Before Tolerance (Required)	The amount of time (in seconds) to allow for clock skew between servers. Valid range is 0 to 3600. The default value is 0.
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 (<code>subject</code>) 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. We recommend to update your applications with the latest version of the agent. We recommend not to change the value of this flag.

5. On the **Actions** screen, click **Download** under **Action Invocation Link**, and then 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. Refer to the documentation of your respective integration kit for more information.


6. Optional: On the **Extended Contract** screen, configure additional attributes for this adapter instance.
7. On the **Summary** screen, review your configuration, modify as needed, and click **Done** to exit the **Create Adapter Instance** workflow.
8. On the **Manage SP Adapter Instances** screen, click **Save** to retain the configuration of the adapter instance.

If you want to exit without saving the configuration, click **Cancel**.

Composite Adapter

For an IdP, PingFederate® includes a Composite Adapter, which 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 may be used to choose an adapter instance based on the method by which a user authenticated, or to integrate an organization's multifactor authentication requirement.


-  **Tip:** For complex authentication requirements, consider implementing one or more authentication policies on the **Policies** screen from the **IdP Configuration** or **SP Configuration** menu.

Configure the Composite Adapter

1. Click **IdP Configuration > Adapters** to open the **Manage IdP Adapter Instances** screen.
2. On the **Manage IdP Adapter Instances** screen, click **Create New Instance** to start the **Create Adapter Instance** workflow.
3. On the **Type** screen, configure the basics of this adapter instance.
 - a) Enter the required information and select the adapter type from the 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 screens.


4. On the **IdP Adapter** screen, configure your Composite Adapter instance as follows:
 - a) Click **Add a new row to 'Adapters'**.
 - b) Select an IdP adapter instance from the **Adapter Instance** list, configure the selected adapter instance (as described in the following table), and then click **Update**.

Column	Description
Policy (Required)	<p>Required (the default) indicates authentication via 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.</p> <p>Sufficient indicates that authentication via 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 <i>not</i> invoked.</p> <p> Important: For the sufficient policy to work correctly, the adapter must return control to PingFederate® after any kind of a failure.</p>
AuthN Context Weight (Required)	<p>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, <i>unless</i> the value is overridden under AuthN Context Override.</p> <p>If weights are the same for two or more contexts, the first one processed is included in the assertion.</p> <p>The default value is 3.</p>
AuthN Context Override	<p>If provided, this value overrides the authentication context value that may be returned from the adapter instance. The value in this field may be sent in the assertion if the associated adapter instance is invoked.</p> <p>If weights are the same for two or more contexts, the first one processed is included in the assertion.</p>

- 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:** Several types of adapters—for example, the IWA IdP Adapter—may be configured to direct end users to an error page if authentication fails for any reason, which will halt further progress through a composite-adapter chain. Ensure that for such adapter instances the **Error URL** option is *not* used in the instance configuration, if continuation through an adapter-chaining sequence is required.

- d) Optional: Use the workflow under **Action** to 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 screen. 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. Select the HTML Form Adapter instance from the **Target Adapter** list.
3. Select a source attribute from the **User ID Selection** list.
4. Click **Update**.


 **Note:** For OAuth use cases, entries in the **Input User Id Mapping** section could override the `login_hint` parameter value provided by the OAuth clients when they submit their requests to the `/as/authorization.oauth2` authorization endpoint.

 **Tip:** 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'** and 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) Defines 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 (**Add to Back**) the 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 may 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** screen, **Add** attributes to be returned from each adapter instance configured on the previous screen.


 **Note:** Attributes must correspond exactly to any or all of the attribute names listed on the Adapter Attribute screens for each configured adapter instance.

6. On the **Adapter Attributes** screen, configure the pseudonym and masking options.

 **Note:** The **Override Attributes** check box in this screen reflects the status of the override option in the **Extended Contract** screen.

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 SP partners use pseudonyms for account linking.

 **Note:** A selection is required regardless of whether 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 (for example, email) to prevent the same pseudonym from being assigned 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** screen, configure the adapter contract for this instance with the following optional workflows:
 - Configure one or more data sources for data store queries.
 - Fulfill adapter contract with values from the adapter (the default), data store queries (if configured), context of the request, text, or expressions (if enabled).
 - Set up the token authorization workflow to validate one or more criteria prior to the issuance of the adapter contract.
 8. On the **Summary** screen, review your configuration, modify as needed, and click **Done** to exit the **Create Adapter Instance** workflow.
 9. On the **Manage IdP Adapter Instances** screen, click **Save** to retain the configuration of the adapter instance. If you want to exit without saving the configuration, click **Cancel**.

Application endpoints

These endpoints provide a means, via standard HTTP, by which external applications can communicate with the PingFederate server.

The SSO and SLO endpoints for an IdP and an 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 (see [Customizable user-facing screens](#) on page 126).

SP endpoints also include those available for SCIM Inbound Provisioning (see [Provisioning for SPs](#) on page 82).

For either SP or IdP servers, a maintenance endpoint is also provided for administrators to verify that the server is running. Endpoints applicable to both server roles also include those needed for adapter-to-adapter mapping (see [Adapter-to-adapter mappings](#) on page 420) and retrieval of WS-Trust metadata (see [WSC and WSP support](#) on page 59).


PingFederate provides a favorite icon for all Application Endpoints. For more information, see [Customize the favicon for application and protocol endpoints](#) on page 144.

IdP endpoints

The following sections describe PingFederate IdP endpoints, including the 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 the parameter TargetResource (or TARGET) is used and includes its own query parameters, the parameter value must be URL-encoded.

Any other parameters that contain restricted characters (many SAML URNs, for example) also must be URL-encoded.

For information about URL encoding, please refer to third party resources, such as [HTML URL-encoding Reference](#) (www.w3schools.com/tags/ref_urlencode.asp).



Parameters are case-sensitive.

/ldp/startSSO.ping

This is the path used to initiate an unsolicited IdP-initiated SSO transaction during which a SAML response containing an assertion is sent to an 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](#) on page 477.

The following table shows the HTTP parameters for this endpoint.

Parameter	Description
PartnerSpId or PARTNER	The federation ID of the SP to whom the SAML response containing an assertion should be issued. This ID is case-sensitive. One of these parameters is required unless the federation ID can be derived from TargetResource or TARGET.
TargetResource or 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. Note that 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 (see Customizable user-facing screens on page 126).
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: <code>urn:oasis:names:tc:SAML:2.0:bindings:HTTP-Artifact</code> <code>urn:oasis:names:tc:SAML:2.0:bindings:HTTP-POST</code> When the parameter is not used, the default ACS URL configured for the SP-partner connection is used, unless an ACS index is specified (see the next parameter, ACSIdx).
ACSIdx (optional - SAML 2.0)	Specifies the index number of partner's ACS (see Set Assertion Consumer Service URLs (SAML) on page 302). 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 may be overridden by policy based on authentication policies. For example, an CIDR Authentication Selector instance could enforce the use of a given adapter instance based on whether a user is on or off the network (see Manage authentication selectors on page 209).
ChangePassword	If a request includes this parameter with a value of <code>true</code> and invokes an HTML Form Adapter instance, the user is redirected to the Change Password form and prompted to update the network password.  Note: In order 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 Configure the HTML Form Adapter on page 431.

Parameter	Description
RequestedFormat (optional - SAML 2.0)	Allows control over the NameId format.
vsid (optional)	Specify the virtual server ID. When absent, PingFederate uses the default virtual server ID (if specified) for the connection (see Identify the SP on page 288) or the SAML federation ID defined in Server Settings (see Specify federation information on page 152).

/idp/startSLO.ping

This is the path used to initiate an IdP-initiated SLO (under SAML 2.0) or an OpenID Connect logout (see [Asynchronous Front-Channel Logout](#) on page 273). 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.

The following table shows the HTTP parameters for this endpoint.

Parameter	Description
TargetResource (optional)	Indicates where the user is redirected after a successful SLO. If this parameter is not included in the request, PingFederate uses as a default the URL for a successful SLO as entered on the IdP Default URL screen. Note that the parameter value must be URL-encoded.
InErrorResource (optional)	Indicates where the user is redirected after an unsuccessful SLO. If this parameter is not included in the request, PingFederate redirects the user to the SLO error landing page hosted within PingFederate (see Customizable user-facing screens on page 126).
Binding (optional - SAML 2.0)	Indicates the binding to be used; 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 SP-partner connection is used (see Specify SLO service URLs (SAML 2.0) on page 304).

/idp/writcdc.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 (optional)	Indicates where the user is redirected after successful IdP Discovery. If this 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. Note that the parameter value must be URL-encoded.

System-services endpoints

These endpoints are the same for both IdPs and SPs and are described under [System-services endpoints](#) on page 460.

SP endpoints

The following sections describe the PingFederate SP endpoints, including the query parameters that each accepts or requires:


- [SP services](#) on page 451
- [SCIM inbound provisioning endpoints](#) on page 455

SP services

The following sections describe PingFederate SP endpoints, including the 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 SP server, for example:

```
https://www.example.com:9031/sp/startSSO.ping
```

 **Important:** When the parameter TargetResource (or TARGET) is used and includes its own query parameters, the parameter value must be URL-encoded.

Any other parameters that contain restricted characters (many SAML URNs, for example) also must be URL-encoded.

For information about URL encoding, please refer to third party resources, such as [HTML URL-encoding Reference](#) (www.w3schools.com/tags/ref_urlencode.asp).


Parameters are case-sensitive.

/sp/startSSO.ping

This is the path used to initiate SP-initiated SSO. In this scenario, the SP issues an SSO request to the IdP asking for an SSO authentication response. Typically, a systems integrator or developer creates one or more links to this endpoint in SP applications to allow users to access various protected resources via 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](#) on page 477.

The following table shows the HTTP parameters for this endpoint.

 **Note:** 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	<p>The federation ID of the IdP that authenticates the user and issues an assertion. This ID is case-sensitive.</p> <p>Required if more than one IdP connection is configured and Domain is not used, and SP authentication policies are turned off.</p> <p>Not required if SP authentication policies are turned on.</p>
SpSessionAuthnAdapterId	<p>The explicit SP adapter instance ID indicating the adapter to use to create an authenticated session or security context.</p> <p>Optional if SP authentication policies are turned off.</p>

Parameter	Description
	Required if SP authentication policies are turned on, unless the PingFederate SP server is able to determine the applicable SP adapter instance based on the target URL mapping configuration and the <code>TargetResource</code> or <code>TARGET</code> value at runtime.
<code>TargetResource</code> or <code>TARGET</code>	This parameter indicates where the end-user is redirected after a successful SSO (the target applications). Note that the parameter value must be URL-encoded. When this parameter is not provided in the URL, a default target resource may be specified in the administrative console, either for all IdP connections (see Configure default URLs on page 336) or for individual connections (see Configure default target URLs on page 363), or both.
<code>InErrorResource</code> (optional)	This parameter indicates where the end-user is redirected after an unsuccessful SSO. If this parameter is not included in the request, PingFederate redirects the user to the SLO error landing page hosted within PingFederate (see Customizable user-facing screens on page 126).
<code>Binding</code> (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: 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 (see Specify SSO service URLs (SAML) on page 357).
<code>AllowCreate</code> (optional - SAML 2.0)	Controls the value of the <code>AllowCreate</code> attribute of the <code>NameIDPolicy</code> element in the <code>AuthnRequest</code> . (The default is <code>true</code> .)
<code>AuthenticatingIdpId</code> (optional - SAML 2.0)	This parameter indicates the preferred IdP for authenticating the user through an IdP proxy, such as PingOne. The parameter specifies the value of the <code>ProviderID</code> attribute in the <code>Scoping/IDPList/IDPEntry</code> element in the <code>AuthnRequest</code> (see section 3.4.1.3.1 of the OASIS SAML document saml-core-2.0-os.pdf). Multiple values are permitted in order to build a preferred list.
<code>ForceAuthn</code> (optional - SAML 2.0 or OpenID Connect)	For SAML 2.0, this parameter controls the attribute of the same name in the <code>AuthnRequest</code> . For OpenID Connect, a value of <code>true</code> sets the prompt parameter in the authentication request to <code>login</code> . For more information about the authentication request and its parameter, please refer to the OpenID Connect specification (openid.net/specs/openid-connect-core-1_0.html#AuthRequest). The default is <code>false</code> .
<code>IsPassive</code> (optional - SAML 2.0 or OpenID Connect)	For SAML 2.0, this parameter controls the attribute of the same name in the <code>AuthnRequest</code> . For OpenID Connect, a value of <code>true</code> sets the prompt parameter in the authentication request to <code>none</code> . The default is <code>false</code> .
<code>RequestedACSIIdx</code>	The index number of your site's Assertion Consumer Service, where you want the assertion to be sent.

Parameter	Description
(optional - SAML 2.0)	
RequestedAcsUrl (optional - SAML 2.0)	The URL of your site's Assertion Consumer Service, where you want the assertion to be sent.
RequestedAuthnCtx (optional - SAML 2.0 or OpenID Connect)	For SAML 2.0, this parameter indicates the requested authentication context of the assertion; allowed values include URIs defined in the SAML specifications (see the OASIS SAML document saml-authn-context-2.0-os.pdf). For OpenID Connect, the specified value becomes the <code>acr_values</code> parameter value in the authentication request. Multiple values are permitted in order to build a preferred list.
RequestedAuthnDeclReq (optional - SAML 2.0)	An alternative to RequestedAuthnCtx, above, indicating the requested authentication context of the assertion by declaring any URI reference (see section 2.7.2.2 of the OASIS SAML document saml-core-2.0-os.pdf). Multiple values are permitted in order to build a preferred list.
RequestedBinding (optional - SAML 2.0)	Indicates the binding requested for the response containing the assertion; allowed values are URIs defined in the SAML specifications.
RequestedFormat (optional - SAML 2.0)	Specifies the value for the Format attribute in the NameIDPolicy element of the AuthnRequest. If not specified, the attribute is not included in the AuthnRequest.
RequestedSPNameQualifier (optional - SAML 2.0)	Indicates that the IdP should return the given name qualifier as part of the assertion (used primarily to identify SP affiliations—see Define SP affiliations on page 327).
vsid (optional)	Specify the virtual server ID. When absent, PingFederate uses the default virtual server ID (if specified) for the connection (see Identify the partner on page 344) or the SAML federation ID defined in Server Settings (see Specify federation information on page 152).
Domain	The domain name associated with the requesting user's IdP to invoke Auto-Connect. In this case, PartnerIdpId cannot be used (see Using Auto-Connect on page 79).

If an adapter is specified in `SpSessionAuthnAdapterId`, 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 screen (see [Configure target URL mapping](#) on page 332).

Note that adapter selection for SP-initiated SSO is similar to that for IdP-initiated SSO except that, because the adapter ID is dependent 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 the user is redirected after a successful SLO. If this parameter is not included in the request, PingFederate uses as a default the URL for a successful SLO, as entered on the SP Default URLs screen. Note that the parameter value must be URL-encoded.
Binding (optional - SAML 2.0)	Indicates the binding to be used; 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 (see Specify SLO service URLs (SAML 2.0) on page 304).
InErrorResource (optional)	Indicates where the user is redirected after an unsuccessful SLO. If this parameter is not included in the request, PingFederate redirects the user to the SLO error landing page hosted within PingFederate (see Customizable user-facing screens on page 126).

/sp/defederate.ping

This is the path used to terminate an account link created during SSO. Account linking provides a means for subject identification on the SP side. Links are created and terminated entirely by a user on the SP side. 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 via account linking on the same PingFederate session, each of those links will be terminated by this endpoint. A local logout is also performed for any link that is terminated.


/sp/cdcstartSSO.ping


This endpoint is used for IdP-Discovery implementations (see [IdP Discovery](#) in the “Supported Standards” chapter of the *Get started with PingFederate* guide). This endpoint is similar to `/sp/startSSO.ping` and accepts the same parameters, with the exception of `PartnerIdpId` and `vsid` (see [/sp/startSSO.ping](#) on page 451). Instead of this parameter, the server attempts to use the common domain cookie to determine the IdP.



/sp/startAttributeQuery.ping

This endpoint is used to initiate an Attribute Query with a SAML 2.0 IdP (see [Attribute Query and XASP](#) in the “Supported Standards” chapter of the *Get started with PingFederate* guide).

The following table shows the HTTP parameters for this endpoint.

 **Note:** Some parameters described below can have multiple values. Specify these values by using multiple independent query string parameters of the same name.

Parameter	Description
Subject	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.
Issuer (optional)	The IssuerDN from the user's X.509 certificate (when 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.

Parameter	Description
PartnerIdpId (except for XASP)	Used to identify the specific IdP partner to which the Attribute Query should be sent. If this parameter is not present, the Subject and Issuer are used to determine the correct IdP.  Note: For XASP, this parameter overrides both the Subject and Issuer parameters.
Format (required for XASP, otherwise optional)	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 (see section 8.3 of the SAML specifications (docs.oasis-open.org/security/saml/v2.0/saml-core-2.0-os.pdf)).  Note: For XASP, this parameter must be set to <code>urn:oasis:names:tc:SAML:1.1:nameid-format:X509SubjectName</code> . If not specified, the parameter defaults to <code>urn:oasis:names:tc:SAML:1.1:nameid-format:unspecified</code> . Note that the parameter must be URL-encoded.
AppId	The unique identifier of the initiating application.
SharedSecret	Used to authenticate the initiating application. The AppId and SharedSecret must both match the application authentication settings within the PingFederate server.
RequestedAttrName (optional)	A name of a user attribute requested from the IdP. For each such desired user attribute, include this parameter. If this parameter is not present, then all allowable user attributes are returned from the IdP. Multiple values are permitted in order to build a preferred list.
vsid (optional)	Specify the virtual server ID. When absent, PingFederate uses the default virtual server ID (if specified) for the connection (see Identify the SP on page 288) or the SAML federation ID defined in Server Settings (see Specify federation information on page 152).

SCIM inbound provisioning endpoints

PingFederate® supports SCIM inbound provisioning and provides four endpoints:

- `/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](http://www.simplecloud.info/specs/draft-scim-core-schema-01.html) (www.simplecloud.info/specs/draft-scim-core-schema-01.html)
- [SCIM Specification](http://www.simplecloud.info/specs/draft-scim-api-01.html) (www.simplecloud.info/specs/draft-scim-api-01.html)

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	<p><code>/pf-scim/v1/Users</code></p> <ul style="list-style-type: none"> 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 data store. The user ID is set as the <code>id</code> attribute in the JSON response, and the full URL to reference the user is in the HTTP response Location header. <p>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 <code>/pf-scim/v1/Users/user_id</code>) and setting the request header <code>X-HTTP-Method-Override</code> value to <code>PUT</code> or <code>DELETE</code>, respectively. (For more information, see the <code>PUT</code> and <code>DELETE</code> method descriptions at the end of this topic.)</p>
GET	<p><code>/pf-scim/v1/Users</code></p> <ul style="list-style-type: none"> Retrieves all user records. A successful response is indicated by an HTTP 200 status code and a list of all users and their attributes. <p><code>/pf-scim/v1/Users/user_id</code></p> <ul style="list-style-type: none"> Retrieve the user record of a specific user. A successful response is indicated by an HTTP 200 status code and the requested user record. <p><code>/pf-scim/v1/Users?attributes=attribute</code></p> <ul style="list-style-type: none"> 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. <p> Note: For more information, see 3.2.2 List/Query Resources in SCIM Specification (www.simplecloud.info/specs/draft-scim-api-01.html#query-resources).</p> <p><code>/pf-scim/v1/Users?filter=filter</code></p> <ul style="list-style-type: none"> 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. <p> Note: For more information, see 3.2.2.1 Filtering in SCIM Specification (www.simplecloud.info/specs/draft-scim-api-01.html#rfc.section.3.2.2.1).</p> <p><code>/pf-scim/v1/Users?sortBy=attribute&sortOrder=ascending descending</code></p> <ul style="list-style-type: none"> 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. Note that, depending on the implementation of the target data store, the target data store may not return the user records that do not contain a value for that specific attribute (as indicated by the <code>sortBy</code> parameter in the request). <p> Note: For more information, see 3.2.2.2 Sorting in SCIM Specification (www.simplecloud.info/specs/draft-scim-api-01.html#rfc.section.3.2.2.2).</p>


HTTP method	Description
	<p><code>/pf-scim/v1/Users?startIndex=x[&count=y]</code></p> <ul style="list-style-type: none"> Retrieves the user records starting with a specific index number, a positive integer <i>x</i>. If the optional count parameter is included (with a positive integer <i>y</i>), 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. <p> Note: For more information, see 3.2.2.3 Pagination in SCIM Specification (www.simplecloud.info/specs/draft-scim-api-01.html#rfc.section.3.2.2.3).</p> <p> Tip: You may use a combination of the aforementioned parameters in one query to narrow your search results.</p>
PUT	<p><code>/pf-scim/v1/Users/user_id</code></p> <ul style="list-style-type: none"> 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 data store. A successful PUT operation returns an HTTP 200 status code and the entire updated user record within the response body.
DELETE	<p><code>/pf-scim/v1/Users/user_id</code></p> <ul style="list-style-type: none"> 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 screen in the applicable IdP connection. A successful response is indicated by an HTTP 200 status code. <p> Note: For a list of HTTP error codes that may be returned, see 3.9 HTTP Response Codes in SCIM Specification (www.simplecloud.info/specs/draft-scim-api-01.html#anchor6).</p>

/pf-scim/v1/Groups






The groups endpoint is where client applications make HTTP requests to create, retrieve, update, and delete groups.

 **Note:** 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** screen 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	<p><code>/pf-scim/v1/Groups</code></p> <ul style="list-style-type: none"> 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 data store. The group ID is set as the <code>id</code> attribute in the JSON response, and the full URL to reference the group is in the HTTP response Location header. <p>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 <code>/pf-scim/v1/Groups/group_id</code>) and setting the request header <code>X-HTTP-Method-Override</code></p>


HTTP method	Description
GET	<p>value to PUT or DELETE, respectively. (For more information, see the PUT and DELETE method descriptions at the end of this topic.)</p> <p><code>/pf-scim/v1/Groups</code></p> <ul style="list-style-type: none"> Retrieves all group records. A successful response is indicated by an HTTP 200 status code and a list of all groups and their attributes. <p><code>/pf-scim/v1/Groups/group_id</code></p> <ul style="list-style-type: none"> Retrieve the group record of a specific group. A successful response is indicated by an HTTP 200 status code and the requested group record. <p><code>/pf-scim/v1/Groups?attributes=attribute</code></p> <ul style="list-style-type: none"> 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. <p> Note: For more information, see 3.2.2 List/Query Resources in SCIM Specification (www.simplecloud.info/specs/draft-scim-api-01.html#query-resources).</p> <p><code>/pf-scim/v1/Groups?filter=filter</code></p> <ul style="list-style-type: none"> 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. <p> Note: For more information, see 3.2.2.1 Filtering in SCIM Specification (www.simplecloud.info/specs/draft-scim-api-01.html#rfc.section.3.2.2.1).</p> <p><code>/pf-scim/v1/Groups?sortBy=attribute&sortOrder=ascending descending</code></p> <ul style="list-style-type: none"> 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. Note that, depending on the implementation of the target data store, the target data store may not return the group records that do not contain a value for that specific attribute (as indicated by the sortBy parameter in the request). <p> Note: For more information, see 3.2.2.2 Sorting in SCIM Specification (www.simplecloud.info/specs/draft-scim-api-01.html#rfc.section.3.2.2.2).</p> <p><code>/pf-scim/v1/Groups?startIndex=x[&count=y]</code></p> <ul style="list-style-type: none"> Retrieves the group records starting with a specific index number, a positive integer <i>x</i>. If the optional count parameter is included (with a positive integer <i>y</i>), 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. <p> Note: For more information, see 3.2.2.3 Pagination in SCIM Specification (www.simplecloud.info/specs/draft-scim-api-01.html#rfc.section.3.2.2.3).</p> <p> Tip: You may use a combination of the aforementioned parameters in one query to narrow your search results.</p>

HTTP method	Description
PUT	<p><code>/pf-scim/v1/Groups/group_id</code></p> <ul style="list-style-type: none"> Updates group attributes for the specified group, using JSON in the body of the HTTP request. Attributes not included in the request are set to a default value in the data store. A successful PUT operation returns an HTTP 200 status code and the entire updated group record within the response body.
DELETE	<p><code>/pf-scim/v1/Groups/group_id</code></p> <ul style="list-style-type: none"> Deletes the group record for the specified group. A successful response is indicated by an HTTP 200 status code.

 **Note:** For a list of HTTP error codes that may be returned, see [3.9 HTTP Response Codes](#) in SCIM Specification (www.simplecloud.info/specs/draft-scim-api-01.html#anchor6).

`/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.

 **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
GET	<p>Retrieves the resource's schema for an IdP connection based on the authentication information.</p> <p>A successful response is indicated by an HTTP 200 status code and the results in the message body.</p>

Sample output

```
$ 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 SP configuration endpoint is where developers can retrieve detailed information on the PingFederate® SCIM 1.1 implementation. When inbound provisioning is enabled for an SP PingFederate server, an HTTP GET request to this endpoint returns a JSON response outlining SCIM 1.1 compliance details.



Note: The /pf-scim/v1/ServiceProviderConfigs endpoint does not require authentication. JSON is currently the only supported format for the HTTP message body.

Sample output

```

$ 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


These endpoints apply to the PingFederate server generally, whether used as an IdP, SP, or both.



Note: Parameters are case-sensitive.

/pf/heartbeat.ping

This endpoint returns an "OK" browser message and an HTTP 200 status indication if the PingFederate runtime server is up and functional. You can customize the message by modifying a PingFederate property and a Velocity template file (see [Customize the heartbeat message](#) on page 144).

 **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.

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 (see [Configure runtime reporting](#) on page 147).


/pf/adapter2adapter.ping

This endpoint initiates direct IdP-to-SP adapter mapping, when that feature is configured (see [Adapter-to-adapter mappings](#) on page 420).

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 uses as a default the URL for a successful SSO (see Configure default URLs on page 336).
InErrorResource (optional)	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 (see Customizable user-facing screens on page 126).
SpSessionAuthnAdapterId (optional)	Indicates the SP adapter instance ID to be used. If not provided and more than one SP adapter instance is configured with adapter-to-adapter mapping, PingFederate uses configured defaults (see Configure target URL mapping on page 332).
IdpAdapterId (optional)	Indicates the IdP adapter to use for authentication if more than one IdP adapter is configured in adapter-to-adapter mappings.
ChangePassword	If a request includes this parameter with a value of <code>true</code> and invokes an HTML Form Adapter instance, the user is redirected to the Change Password form and prompted to update the network password.  Note: In order 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 Configure the HTML Form Adapter on page 431.

/pf/sts.wst

This endpoint initiates direct STS token-to-token exchange and token validation from an IdP token processor to an SP token generator, when that feature is configured (see [Token translator mappings](#) on page 423).

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, a secondary PingFederate listening port must be configured and used by partners or STS clients for the relevant endpoints—*.ssaml* and *.wst (see [Configure PingFederate properties](#) on page 93).

/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.
vsid (optional)	Specify the virtual server ID. If absent, PingFederate uses the default virtual server ID (if specified) for the connection or the federation ID defined in Server Settings (see Specify federation information on page 152).

/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 (optional)	Specify the virtual server ID. 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 in the Server Configuration > Server Settings > Federation Info screen.

 **Note:** 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.

Construct an alternative metadata exchange endpoint

You can embed virtual server ID information into an STS metadata exchange endpoint or a SAML and WS-Federation metadata exchange endpoint. This is useful for scenarios where partners prefer to retrieve metadata

by sending one query parameter (`PartnerSpId` or `PartnerIdpId`) instead of two query parameters (`PartnerSpId` or `PartnerIdpId` *and* `vsid`).

1. Construct a JSON object containing a key-value pair of the virtual server ID by using the following format:

```
{"vsid": "<VirtualServerIdValue>"}
```

For example, if the virtual server ID is `Engineering`, the JSON object is:

```
{"vsid": "Engineering"}
```

2. Base64url-encode the JSON object.

For example, if the JSON object is `{"vsid": "Engineering"}`, the base64url-encoded value is:

```
eyJ2c2lkIjoiRW5naW5lZXJpbmcfQ
```

For more information about base64url, see [RFC4648](https://tools.ietf.org/html/rfc4648) (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/sts_mex.ping`)

Between `/pf` and `/federation_metadata.ping`

STS metadata endpoint (`/pf/sts_mex.ping`)

Between `/pf` and `/sts_mex.ping`

for example, if the base64url-encoded value is `eyJ2c2lkIjoiRW5naW5lZXJpbmcfQ`, the metadata exchange endpoints embedding with the virtual server ID are:

Federation metadata endpoint

```
/pf/eyJ2c2lkIjoiRW5naW5lZXJpbmcfQ/federation_metadata.ping
```

Example: `https://idp.example.com:9031/pf/eyJ2c2lkIjoiRW5naW5lZXJpbmcfQ/federation_metadata.ping?PartnerSpId=sp.example.org`

STS metadata endpoint

```
/pf/eyJ2c2lkIjoiRW5naW5lZXJpbmcfQ/sts_mex.ping
```

Example: `https://idp.example.com:9031/pf/eyJ2c2lkIjoiRW5naW5lZXJpbmcfQ/sts_mex.ping?PartnerSpId=sp.example.org`

OAuth 2.0 endpoints

When developing OAuth-capable applications, developers must follow [the OAuth 2.0 Authorization Framework](https://tools.ietf.org/html/rfc6749) (tools.ietf.org/html/rfc6749) and, if applicable, [the OpenID Connect specification](https://openid.net/connect) (openid.net/connect), which means that the applications must send requests to various OAuth endpoints to obtain authorization grants, access tokens, and (if applicable) refresh tokens and ID tokens. Furthermore, there are also endpoints for the clients to revoke tokens, the resource owners to revoke authorization grants, and the clients to retrieve OpenID Connect metadata, such as the User Info endpoint.

The topic describes OAuth-developer information on PingFederate® endpoints for the OAuth AS. Unless otherwise indicated, these endpoints and associated parameters are defined in the aforementioned specifications



Note: Begin each URL with the fully qualified server name and port number of your IdP or SP PingFederate server; for example:

- `https://www.example.com:9031/as/authorization.oauth2`
- `https://www.example.com:9031/as/token.oauth2`

Authorization endpoint

The authorization endpoint is defined in [the OAuth 2.0 Authorization Framework](https://tools.ietf.org/html/rfc6749#section-3.1) (tools.ietf.org/html/rfc6749#section-3.1) and is used by the OAuth AS to interact directly with resource owners, authenticate them, and obtain their authorization. Typically, an OAuth client makes an authorization request by directing a resource owner, via an HTTP user-agent, to the authorization endpoint. After completing 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.








Note: This endpoint may be used as part of 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.

Endpoint: `/as/authorization.oauth2`

The table below shows parameters for this endpoint:

Parameter	Description
<code>client_id</code> (Required)	The client identifier.
<code>response_mode</code>	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 (openid.net/specs/oauth-v2-form-post-response-mode-1_0.html).
<code>response_type</code> (Required)	A value of <code>code</code> results in the Authorization Code grant type while a value of <code>token</code> implies the Implicit grant type. Additionally, a value of <code>id_token</code> can be requested by implicit clients.
<code>code_challenge</code>	Applicable only when <code>response_type</code> is <code>code</code> . Supply a one-time string value used to associate the authorization request with the token request to reduce the risk of code interception attack. For more information, refer to Proof Key for Code Exchange (PKCE) by OAuth Public Clients (tools.ietf.org/html/rfc7636). Note: If used, the OAuth client must submit the corresponding code verifier when using the authorization code to obtain an access token (see <code>code_verifier</code> in Authorization code grant type on page 469).
<code>code_challenge_method</code>	Applicable only when the <code>response_type</code> parameter value is <code>code</code> and a <code>code_challenge</code> parameter value is provided. This parameter indicates the transformation method used to derive the <code>code_challenge</code> parameter value from that of the <code>code_verifier</code> parameter. PingFederate® OAuth AS supports two transformation methods: <ul style="list-style-type: none"> <code>plain</code>, which indicates the <code>code_challenge</code> parameter value is that of the <code>code_verifier</code> parameter. <code>S256</code>, which indicates the <code>code_challenge</code> parameter is derived from the <code>code_verifier</code> parameter value as follows: <pre>code_challenge=Base64Url-encode(SHA256(ASCII(code_verifier))), where:</pre> <ul style="list-style-type: none"> <code>ASCII(code_verifier)</code> denotes the octets of the ASCII representation of the <code>code_verifier</code> value. <code>SHA256(octets)</code> denotes the SHA 256-bit hash of the octets. <code>Base64Url-encode(octets)</code> denotes the base64url encoding of octets; the output is URL-safe.

Parameter	Description
	<p> Note: For detailed information about the transformation method, refer to <i>Proof Key for Code Exchange (PKCE) by OAuth Public Clients</i> (tools.ietf.org/html/rfc7636).</p> <p>The code_challenge_method parameter value is case-sensitive. An error message is returned to the clients for any other values.</p> <p>Note that omitting the code_challenge_method parameter has the same effect as providing the code_challenge_method parameter with a value of plain.</p>
redirect_uri	<p>Required if more than one redirection URIs are configured in PingFederate for the client or if a wildcard is used for a single URI entry. Optional for clients with only one specific redirection URI configured.</p> <p>Note that if this parameter is used, the same parameter and value must also be used in subsequent token requests (see <i>OAuth grant type parameters</i> on page 469).</p>
claims_locales	<p>Specifies the end-user's preferred languages for claims being returned in a space-separated list, ordered by preference. The values must conform to guidelines defined under <i>IETF BCP 47</i> (tools.ietf.org/html/bcp47).</p> <p> 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 screen in the IdP Adapter Mapping configuration or the OAuth Attribute Mapping configuration in an IdP connection).</p>
login_hint	<p>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.</p>
max_age	<p>The allowable elapsed time (in seconds) since the end users last authenticated. If the elapsed time exceeds the value of max_age, the end users are prompted for reauthentication.</p> <p> Tip: The HTML Form Adapter supports the max_age parameter by tracking the authentication time for each user.</p>
scope	<p>The scope of the access request expressed as a list of space-separated, case-sensitive strings. Valid scope values are defined in the OAuth Settings > Authorization Server Settings screen.</p> <p> Tip: Scopes can be restricted per client.</p>
state	<p>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.</p>
ui_locales	<p>Specifies the end-user's preferred languages for OAuth user interactions in a space-separated list, ordered by preference. The values must conform to guidelines defined under <i>IETF BCP 47</i> (tools.ietf.org/html/bcp47).</p>
idp (or PartnerIdpId)	<p>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.</p>
pfdpadapterid (or IdpAdapterId)	<p>A PingFederate OAuth AS parameter indicating the IdP adapter instance ID of the adapter to use for user authentication.</p>

Parameter	Description
	 Note: This parameter may 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, users are presented 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 is displayed on the user authorization page.


OpenID Connect parameters

The table below displays OpenID Connect parameters for this endpoint:

Parameter	Description
<code>acr_values</code>	Specifies the Authentication Context Class Reference (<code>acr</code>) values for the AS to use when processing an Authentication Request. Express as a space-separated string, listing the values in order of preference.
<code>id_token_hint</code>	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.
<code>nonce</code>	Specifies a string value used to associate a client session with an ID token and to reduce replay attacks. The value is passed through unmodified from an authorization request to the ID token.
<code>prompt</code>	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, this parameter can be used by the client to verify that the end user is still present for the current session or to bring attention to the request. PingFederate supports the following values: <code>none</code> , <code>login</code> , <code>consent</code> .

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
<code>access_token_manager_id</code>	<p>The <code>access_token_manager_id</code> 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.</p>  Note: When the <code>access_token_manager_id</code> parameter is specified, PingFederate ignores the <code>aud</code> parameter.
<code>aud</code>	The <code>aud</code> 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 which would be contained within the path of the preconfigured 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 1: 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 2: An exact match is a better match than a partial match

Access Token Management Instances	Resource URIs
ATM1	<code>https://localhost:9031/app1</code> <code>https://localhost:9031/app2/data</code> <code>https://app.example.local</code>
ATM2	<code>https://localhost:9031/app1/data</code> <code>https://localhost:9031/app2/data/get</code>

`https://localhost:9031/app1` (a resource URI preconfigured for ATM1) is a partial match for `https://localhost:9031/app1/data` (the provided URI). However, ATM2 is chosen because `https://localhost:9031/app1/data` (a resource URI preconfigured for ATM2) is an exact match against the provided URI.

Example 3: 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, ATM2 is chosen because `https://localhost:9031/app2/data/get` matches more specifically against the provided URI.

Token endpoint

The token endpoint is defined in the [the OAuth 2.0 Authorization Framework](https://tools.ietf.org/html/rfc6749#section-3.2) (tools.ietf.org/html/rfc6749#section-3.2) and used by the client to obtain an access token and possibly a refresh token by presenting its authorization grant. The token endpoint is used with every authorization grant except for the Implicit grant type (since an access token is issued directly from the authorization endpoint).

Endpoint: `/as/token.oauth2`



Note: Per OAuth specifications, this endpoint accepts only the HTTP POST method.

OAuth client identification and authentication

OAuth clients can authenticate to the OAuth AS using this endpoint by presenting their client identifier and client secret either using the HTTP Basic authentication scheme (where the client identifier is the username, and the client secret is the password) or with the following HTTP request parameters:


Parameter	Description
<code>client_id</code>	The client identifier.
<code>client_secret</code>	The client secret.

Whenever possible, the use of HTTP Basic is recommended over the use of the request parameters.

Clients without a client secret can use the `client_id` parameter to identify themselves to the OAuth AS and omit the `client_secret` parameter.

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
<code>access_token_manager_id</code>	<p>The <code>access_token_manager_id</code> 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.</p> <p> Note: When the <code>access_token_manager_id</code> parameter is specified, PingFederate ignores the <code>aud</code> parameter.</p>
<code>aud</code>	<p>The <code>aud</code> 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.</p>

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 which would be contained within the path of the preconfigured 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 1: 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 2: An exact match is a better match than a partial match

Access Token Management Instances	Resource URIs
ATM1	<p><code>https://localhost:9031/app1</code></p> <p><code>https://localhost:9031/app2/data</code></p> <p><code>https://app.example.local</code></p>
ATM2	<p><code>https://localhost:9031/app1/data</code></p> <p><code>https://localhost:9031/app2/data/get</code></p>


`https://localhost:9031/app1` (a resource URI preconfigured for ATM1) is a partial match for `https://localhost:9031/app1/data` (the provided URI). However, ATM2 is chosen because `https://localhost:9031/app1/data` (a resource URI preconfigured for ATM2) is an exact match against the provided URI.

Example 3: 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, ATM2 is chosen because `https://localhost:9031/app2/data/get` matches more specifically against the provided URI.

OAuth grant type parameters

Other parameters accepted by the `/as/token.oauth2` endpoint vary by the grant type being presented. They also include both OAuth-defined standard parameters and parameters proprietary to PingFederate®. The grant type of the access token request is indicated by the following parameter:

Parameter	Description
grant_type (Required)	<p>Indicates the type of grant being presented in exchange for an access token and possibly a refresh token. The value is an extensibility mechanism of the OAuth 2.0 specification. PingFederate supports these values:</p> <ul style="list-style-type: none"> • <code>authorization_code</code> • <code>refresh_token</code> • <code>password</code> • <code>client_credentials</code> • <code>urn:ietf:params:oauth:grant-type:jwt-bearer</code> • <code>urn:ietf:params:oauth:grant-type:saml2-bearer</code> • <code>urn:pingidentity.com:oauth2:grant_type:validate_bearer</code> <p> Note: Further parameters associated with each grant type are defined in the following sections.</p>

Authorization code grant type

These parameters apply when the `grant_type` parameter for `/as/token.oauth2` 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 <code>response_type</code> parameter is <code>code</code> .
code_verifier	<p>Required if the authorization request was sent with a <code>code_challenge</code> parameter to reduce the risk of code interception attack.</p> <p>Based on the <code>code_challenge_method</code> parameter value (if provided in the request for the authorization code in the first place), PingFederate OAuth AS validates the <code>code_verifier</code> parameter value against that of the <code>code_challenge</code> value. If the validation returns no error, PingFederate OAuth AS returns an access token (provided that there is no other error condition); otherwise it returns an error to the client.</p> <p>For more information about the <code>code_challenge</code> parameter, the <code>code_challenge_method</code> parameter, and the support for <i>Proof Key for Code Exchange (PKCE) by OAuth Public Clients</i> (tools.ietf.org/html/rfc7636), see <i>Authorization endpoint</i> on page 464.</p>
redirect_uri	<p>This parameter is required if the <code>redirect_uri</code> parameter was included in the authorization request that resulted in the issuance of the code (see <i>Authorization endpoint</i> on page 464). The value here must match the authorization-request value, if applicable.</p> <p>The parameter is also required for clients with multiple redirection URIs or one redirection URI that uses wildcards.</p> <p>The parameter is optional for clients with only one specific redirection URI.</p>

Refresh token grant type

These parameters apply when the `grant_type` parameter for `/as/token.oauth2` is set to `refresh_token`.

Parameter	Description
refresh_token (Required)	The refresh token issued to the client during a previous access-token request.
scope	<p>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 by the resource owner. If omitted, the scope is treated as equal to that originally granted by the resource owner.</p> <p>Valid scope values are defined in the OAuth Settings > Authorization Server Settings screen.</p> <p>Scopes can also be restricted per client.</p>

Resource owner password credentials grant type


These parameters apply when the `grant_type` parameter for `/as/token.oauth2` is set to `password`.

Parameter	Description
username (Required)	The username, encoded as UTF-8.
password (Required)	The password, encoded as UTF-8.
scope	The scope of the access request.
validator_id	A PingFederate OAuth AS parameter indicating the instance ID of the password credential validator to be used to check the username and password (and the associated attribute mapping into the <code>USER_KEY</code> of the persistent grant). If multiple validator instances are configured and mapped and no <code>validator_id</code> parameter is provided, each instance will be tried sequentially until one succeeds or they all fail.

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; for 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 detected by the LDAP PCV. OAuth-client developers may create custom experiences based on the error messages.

 **Tip:** These customizable messages are stored in the `<pf_install>/pingfederate/server/default/conf/language-packs/pingfederate-messages.properties` file.

As needed, they may be localized using the PingFederate localization framework for an international audience.

Note that the `client_id` parameter is not required when the **Allow unidentified clients to make resource owner password credentials grants** check box is selected on the **OAuth Settings > Authorization Server Settings** screen.

Client credentials grant type

These parameters apply when the `grant_type` parameter for `/as/token.oauth2` is set to `client_credentials`.

Parameter	Description
<code>scope</code>	The scope of the access request.

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
<code>assertion</code> (Required)	A JSON Web Token (JWT), as defined in RFC7523, section 2.1 (tools.ietf.org/html/rfc7523#section-2.1).
<code>scope</code>	The scope of the access request.

Note that the `client_id` parameter is not required when the **Allow unidentified clients to request extension grants** check box is selected on the **OAuth Settings > Authorization Server Settings** screen.

SAML 2.0 Bearer Assertion grant type

These parameters apply when the `grant_type` parameter for `/as/token.oauth2` is set to `urn:ietf:params:oauth:grant-type:saml2-bearer`.

Parameter	Description
<code>assertion</code> (Required)	A single SAML 2.0 assertion, which must be encoded using <code>base64url</code> , as described in RFC4648, section 5 (tools.ietf.org/html/rfc4648#section-5).
<code>scope</code>	The scope of the access request.


Note that the `client_id` parameter is not required when the **Allow unidentified clients to request extension grants** check box is selected on the **OAuth Settings > Authorization Server Settings** screen.

Access token validation grant type

These parameters apply when the `grant_type` parameter for `/as/token.oauth2` is set to `urn:pingidentity.com:oauth2:grant_type:validate_bearer`.

Parameter	Description
<code>token</code> (Required)	The bearer access token to be validated.

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. The grant type allows an RS to check with the OAuth AS on the validity of a bearer access token that it has received from a client making a protected-resources call.

 **Tip:** Alternatively, an RS client can use the standard-based *Introspection Endpoint* (at `/as/introspect.oauth2`) to validate an access token or a refresh token.

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. The `expires_in` element indicates for how many more seconds the token is valid; note that the value may increase on subsequent validation calls if a token lifetime extension policy is in place (applicable only to access tokens using the reference-token data model).

Example

```
{
  "scope": "read edit admin",
  "token_type": "urn:pingidentity.com:oauth2:validated_token",
  "expires_in": 3172,
  "client_id": "super_cool_mobile_client",
  "access_token":
  {
    "uid": "sfHqhad9onMjXsQNI1mZP9mD7AQasmskd",
    "group": ["employee", "sales", "manager"],
    "email": "someguy@example.cloud"
  }
}
```

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 (see [Configure an OAuth client](#) on page 269).

Introspection endpoint

The introspection endpoint is defined in the [OAuth 2.0 Token Introspection](#) specification (tools.ietf.org/html/rfc7662) and used by a resource server (RS) client to validate an access token or a refresh token prior to granting access to a protected-resources call.

Endpoint: `/as/introspect.oauth2`



Note: Per OAuth specifications, this endpoint accepts only the HTTP POST method.

OAuth client identification and authentication

OAuth clients can authenticate to the OAuth AS using this endpoint by presenting their client identifier and client secret either using the HTTP Basic authentication scheme (where the client identifier is the username, and the client secret is the password) or with the following HTTP request parameters:


Parameter	Description
<code>client_id</code>	The client identifier.
<code>client_secret</code>	The client secret.

Whenever possible, the use of HTTP Basic is recommended over the use of the request parameters.

Clients without a client secret can use the `client_id` parameter to identify themselves to the OAuth AS and omit the `client_secret` parameter.

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
<code>access_token_manager_id</code>	<p>The <code>access_token_manager_id</code> 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.</p> <p> Note: When the <code>access_token_manager_id</code> parameter is specified, PingFederate ignores the <code>aud</code> parameter.</p>
<code>aud</code>	<p>The <code>aud</code> 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.</p>

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 which would be contained within the path of the preconfigured 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 1: 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 2: An exact match is a better match than a partial match

Access Token Management Instances	Resource URIs
ATM1	<p><code>https://localhost:9031/app1</code></p> <p><code>https://localhost:9031/app2/data</code></p> <p><code>https://app.example.local</code></p>
ATM2	<p><code>https://localhost:9031/app1/data</code></p> <p><code>https://localhost:9031/app2/data/get</code></p>

`https://localhost:9031/app1` (a resource URI preconfigured for ATM1) is a partial match for `https://localhost:9031/app1/data` (the provided URI). However, ATM2 is chosen because `https://localhost:9031/app1/data` (a resource URI preconfigured for ATM2) is an exact match against the provided URI.

Example 3: 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, ATM2 is chosen 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 (see [Configure an OAuth client](#) on page 269).

Introspection parameters

The RS acts in the role of a client for the request/response exchange with the PingFederate® OAuth AS to make the validation call using the following parameters:

Parameter	Description
token (Required)	The bearer access token or refresh token to be validated.
token_type_hint	A hint about the type of token submitted for validation. PingFederate supports the following values: <ul style="list-style-type: none"> access_token refresh_token Required only when validating a refresh token.

Client authentication is not required. In other words, 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. Note that 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.
- 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.

If a token is invalid, the OAuth AS returns {"active":false} to the client.

Example 1: A response for a valid access token

```
{
  "sub":"joe",
  "Username":"joe",
  "active":true,
  "OrgName":"Ping Identity Corporation",
  "token_type":"Bearer",
  "exp":1467008819,
  "client_id":"ac_client",
  "username":"joe"}
```

Example 2: A response for a valid refresh token

```
{"active":true}
```


Example 3: A response for an invalid token


```
{"active":false}
```

Token revocation endpoint

The token revocation endpoint is defined in the [OAuth 2.0 Token Revocation](https://tools.ietf.org/html/rfc7009) specification (tools.ietf.org/html/rfc7009). It 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.

Endpoint: /as/revoke_token.oauth2

 **Note:** Per OAuth specifications, this endpoint accepts only the HTTP POST method.

 **Important:** Direct access token revocation is only supported for *Internally Managed Reference Tokens*. Access tokens of the type *JSON Web Token (JWT)* 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.

OAuth client identification and authentication

OAuth clients can authenticate to the OAuth AS using this endpoint by presenting their client identifier and client secret either using the HTTP Basic authentication scheme (where the client identifier is the username, and the client secret is the password) or with the following HTTP request parameters:

Parameter	Description
client_id	The client identifier.
client_secret	The client secret.

Whenever possible, the use of HTTP Basic is recommended over the use of the request parameters.

Clients without a client secret can use the client_id parameter to identify themselves to the OAuth AS and omit the client_secret parameter.

Parameters

The token revocation endpoint uses the following parameters using the `application/x-www-form-urlencoded` format in the HTTP request entity-body:

Parameter	Description
token	The token that the client wants to revoke.
(Required)	
token_type_hint	A hint about the type of token submitted for revocation. PingFederate® supports the following values: <ul style="list-style-type: none"> access_token refresh_token


If a refresh token is revoked, the associated access grant and access tokens will be revoked as well. If an access token is revoked, 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 **OAuth Settings > Authorization Server Settings** screen, the implicit access grant will also be revoked with the access token.

Grant-management endpoint

The grants endpoint (two are provided, one for use with parameters) is where end-users/resource owners go to view (and optionally revoke) the persistent access grants they have made. This endpoint is not part of the OAuth specification, but many OAuth providers offer a similar type of functionality. The grants displayed are those associated with the `USER_KEY` of the authenticated user. The same attribute mapping(s) from the authentication source to `USER_KEY` 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 shows 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 IdP with whom to initiate Browser 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 exactly the same way as the authorization endpoint.

OpenID Provider configuration endpoint

This public endpoint provides configuration information for the OAuth clients to interface with PingFederate® using the OpenID Connect protocol.

Endpoint: `/.well-known/openid-configuration`

The OpenID Provider configuration endpoint supports HTTP GET request without authentication or parameter; for example:

```
$ curl -s https://localhost:9031/.well-known/openid-configuration|python -m
  json.tool
{
  "authorization_endpoint": "https://localhost:9031/as/
authorization.oauth2",
  ...
  "claims_supported": [
    "address",
    ...
  ],
  "id_token_signing_alg_values_supported": [
    "none",
    "HS256",
    ...
  ],
  "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",
  ...
  "revocation_endpoint": "https://localhost:9031/as/revoke_token.oauth2",
  ...
  "token_endpoint": "https://localhost:9031/as/token.oauth2",
  ...
  "userinfo_endpoint": "https://localhost:9031/idp/userinfo.openid",
  ...
}
```



The UserInfo endpoint (`userinfo_endpoint`) is where clients send access tokens to PingFederate in exchange for a list of additional claims about the users.

The JWKS (`jwks_uri`) endpoint returns a set of public keys that are generated and rolled automatically by PingFederate. Clients can use this information to verify the integrity of asymmetrically-signed ID Tokens.

For information about other parameters, see [OpenID Provider Metadata](#) in OpenID Connect Discovery 1.0 (openid.net/specs/openid-connect-discovery-1_0.html#ProviderMetadata).

Web Service interfaces

PingFederate provides two built-in, SOAP-accessible Web Services related to Browser SSO. These services may be used by client applications to manage partner connections and support integration of web applications, respectively:

- Connection Management Service — Enables creation and deletion of single connection configurations in PingFederate. This service may 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.
 -  **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 (see [Automating configuration migration](#) on page 120).
- SSO Directory Service — Provides web application developers with information regarding partner connections and adapter instances.
 -  **Tip:** Applications accessing the Connection Management Service must first authenticate themselves to the PingFederate server. SSO Directory Service authentication is optional by default, but may be required. For more information, see [Authentication](#) on page 195.

PingFederate also provides the following REST-based Web Services:


- [OAuth Client Management Service](#) on page 483 — Useful for managing OAuth client applications, where needed (see [About OAuth](#) on page 61 and [Manage OAuth clients](#) on page 269).
- [OAuth Access Grant Management Service](#) on page 488 — Allows retrieval and revocation of access grants.
- [Session Revocation API endpoint](#) on page 490 — Allows OpenID Connect clients to query revocation status of their sessions and add end-user sessions to the revocation list.
- [PingFederate administrative API](#) on page 492 — Allows developers and non-developers to change PingFederate IdP Connections settings (see [Service provider SSO configuration](#) on page 330).

Connection Management Service

The Connection Management Service supports basic connection management capabilities and is accessible only on a PingFederate server running the administrative console. This feature is useful in a variety of circumstances, but the following primary use cases were considered:

- As a utility to migrate changes to a partner connection through staging environments (for example: development, test, production).
 - Changes to URLs and keys may be needed to make the connection appropriate to the next environment.
- As a way for an external application to update or delete connections programmatically, or create new ones using an exported connection XML file as a template.

The WAR file for this service, `pf-mgmt-ws.war`, is located in the `<pf_install>/pingfederate/server/default/deploy2` directory.

 **Note:** If you do not want to allow use of the service, it should not be deployed: 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:

```
https://<host_server>:<admin_console_port>/pf-mgmt-ws/ws/ConnectionMigrationMgr?wsdl
```

Exporting a connection

You can export a connection either manually, using the administrative console, or programmatically, via a call to the Connection Management Service.

In either case, 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, data stores, and keypairs. Adapter instances and data stores are referenced by ID, and keypairs 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).

Exporting manually

For information about using the administrative console to export SP connections at an IdP site, see [Access SP connections](#) on page 283.

For information about exporting IdP connection at an SP site, see [Access IdP connections](#) on page 340.

Using the connection service

The Connection Web Service exposes the following method for exporting connections:

```
public String getConnection( String entityId, String role, ) throws
    IOException
```


Code sample

```
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 (data stores, keypairs, 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 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.

 **Tip:** Alternatively, you can import XML connection files via the PingFederate administrative console (see [Access SP connections](#) on page 283 or [Access IdP connections](#) on page 340). 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—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

Below is example client code using the Apache AXIS libraries that invokes 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\"
...
</EntityDescriptor>";
boolean allowUpdate = false;
call.invoke(new Object[]{xml, allowUpdate});
```

Deleting 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—IDP or SP.

Code sample

Below is example client code using the Apache AXIS libraries that invokes 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. This allows a client of this Web Service to create or update a new connection (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:

```
https://<host_server>:<admin_console_port>/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

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. Therefore, because the XML is complex and validation is limited, attempting to create an XML connection from scratch is *not recommended*. Rather, the administrative console should be used to create the initial connection. 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 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.) This service allows web applications to avoid storing and maintaining the data locally. These types of data can be retrieved:

- A list of IdP partners
- A list of SP partners
- A list of IdP adapter instances
- A list of SP adapter instances

The SSO Directory Service provides information useful for integrating an application with a PingFederate server. It is a way for the application to find out dynamically which partners can be used for SSO. This means applications need not be modified when new partners are configured in PingFederate.

The WAR file for this module, `pf-ws.war`, is located in the `pingfederate/server/default/deploy` directory.



Note: If you do not want to allow use of the service, it should not be deployed: remove the WAR file from the `deploy` directory.


The service endpoint is `pf-ws/services/SSODirectoryService`.

The WSDL document describing this service can be retrieved from:


```
http(s)://<pf_runtime_host>:<runtime_port>/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> (see [SOAP request and response examples](#) on page 482).

- `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 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 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 (see [Authentication](#) on page 195). We recommend the use of HTTPS 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. For example, the following is a typical SOAP request for an IdP list using the SSO Directory Service.

```

<?xml version="1.0" encoding="UTF-8"?>
<soapenv:Envelope
  xmlns:soapenv="http://schemas.xmlsoap.org/soap/envelope/"
  xmlns:xsd="http://www.w3.org/2001/XMLSchema"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance">
  <soapenv:Body>
    <ns1:getIDPList
      soapenv:encodingStyle=
        "http://schemas.xmlsoap.org/soap/encoding/"
      xmlns:ns1=
        "https://localhost:9031/ssodir/services/
        SSODirectoryService"/>
    </soapenv:Body>
  </soapenv:Envelope>

```

The PingFederate server's Web Service will return 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/
soap/envelope/"
  xmlns:xsd="http://www.w3.org/2001/XMLSchema"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance">
  <soapenv:Body>
    <getIDPListResponse
      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"
      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"
      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 service is provided primarily for organizations with many OAuth clients, allowing programmatic management of OAuth clients, as an alternative to using the administrative console (see [Manage OAuth clients](#) on page 269).



Note: This Web Service is active only if client records are managed in a database, rather than in XML files (the installed default storage method). For information on setting up database management for OAuth clients, see [Define an OAuth client data store](#) on page 171.



Tip: PingFederate administrative API can also be used to manage OAuth clients programmatically regardless if the client records are managed in XML files or in a database. For more information, see [PingFederate administrative API](#) on page 492.

Endpoints

The REST resources in the following sections are URL path extensions of the PingFederate runtime endpoint:

`http(s)://<pf_runtime_host>:<runtime_port>/pf-ws/rest`

For example: `https://my_corp.com:9031/pf-ws/rest/oauth/clients`



Note: POST and PUT methods described in this section require parameter name/value pairs formatted in JavaScript Object Notation (JSON).



Important: Applications must authenticate to the Web Service using HTTP Basic credentials specified in a PingFederate Password Credential Validator (see [Manage password credential validators](#) on page 196). The configured Credential Validator, in turn, must be selected in the OAuth AS configuration (see [Configure authorization server settings](#) on page 243).

/oauth/clients

This resource accepts the methods [POST](#) on page 483, [PUT](#) on page 486 and [GET](#) on page 487.

POST

Description


Creates a new client based on the parameters provided in the request. Parameters correspond to administrative-console fields; for additional information, see [Manage OAuth clients](#) on page 269.

The required MIME type is `application/json`.

JSON Parameters

Parameter	Description
clientId	(Required) A unique ID for the client.
name	(Required) A descriptive name for the client.
refreshRolling	Indicates whether a new refresh token is issued with each new access token (see Configure an OAuth client on page 269). Allowed values: <code>true</code> or <code>false</code> .

Parameter	Description
	When not provided, the global setting for the AS is used (see Configure authorization server settings on page 243).
redirectUri	The URI(s) to which the OAuth AS redirects the resource owner's user agent after authorization is obtained. <i>Required</i> for Implicit and Authorization Code grant types (see <code>grantTypes</code> below).
logoUrl	A URL for the logo presented to the user on the grant revocation page.
secret	Client password or phrase, <i>required</i> for the Client Credentials grant type unless client-certificate authentication is needed (see next two parameters).
clientCertIssuerDn	Client TLS certificate issuer, <i>required</i> for the Client Credentials grant type unless a <code>secret</code> is provided.
clientCertSubjectDn	Client TLS certificate subject, <i>required</i> for the Client Credentials grant type unless a <code>secret</code> is provided.
description	A description of what the client application does, displayed in browser when the user is prompted for authorization.
persistentGrantExpirationType	Indicates whether to override the global setting for the AS (see Configure authorization server settings on page 243). Allowed values: <ul style="list-style-type: none"> • <code>SERVER_DEFAULT</code> (the default) – Use the global setting for the AS. • <code>NONE</code> – Grants do not expire, regardless of the global setting. • <code>OVERRIDE_SERVER_DEFAULT</code> – Use with both of the <code>persistentGrant*</code> parameters below to set the expiration time period.
persistentGrantExpirationTime	An integer representing units of time for storage of persistent grants for this client—use with <code>persistentGrantExpirationTimeUnit</code> (see below) and only when <code>persistentGrantExpirationType</code> is set to <code>OVERRIDE_SERVER_DEFAULT</code> .
persistentGrantExpirationTimeUnit	Units for the expiration time set in the parameter above (if applicable). Allow values: <ul style="list-style-type: none"> • <code>h</code> (hours) • <code>d</code> (days)
bypassApprovalPage	If set to <code>true</code> , user consent to resource access is assumed and the approval page is not presented.
restrictScopes	If set to <code>true</code> , limits client access to a subset of the scopes defined for the AS (see Configure authorization server settings on page 243). Scopes are limited to the default scope and any listed for <code>restrictedScopes</code> (see below). If no scopes are listed and this parameter is <code>true</code> , only the default scope is available for the client.
restrictedScopes	When used with <code>restrictScopes</code> , limits access to the scope(s) provided in the JSON list, in addition to the default scope.
grantTypes	One or more grant types allowed for the client (see Grant types on page 62). Allowed values: <ul style="list-style-type: none"> • <code>authorization_code</code> • <code>password</code> (Resource Owner Password Credentials) • <code>refresh_token</code> (Use with <code>authorization_code</code> or <code>password</code>.) • <code>client_credentials</code> • <code>implicit</code> • <code>extension</code> (SAML 2.0 Bearer) • <code>urn:pingidentity.com:oauth2:grant_type:validate_bearer</code> (Access Token Validation)

Parameter	Description
	 Note: Provide at least one grant type. To provide multiple grant types, include them into a comma-separated list.
defaultAccessTokenManagerId	The default access token manager for this client (see OAuth access token management on page 245).
validateUsingAllEligibleAtms	<p>Applicable only to resource server (RS) clients.</p> <p>If set to <code>true</code>, this RS client is not required to specify additional parameters (<code>access_token_manager_id</code> or <code>aud</code>) to disambiguate the ATM instance in its token validation requests. When the RS 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.</p>
idTokenSigningAlgorithm	<p>The JSON Web Signature (JWS) algorithm required for the ID Token. Allowed values:</p> <ul style="list-style-type: none"> • none - No signing algorithm • HS256 - HMAC using SHA-256 • HS384 - HMAC using SHA-384 • HS512 - HMAC using SHA-512 • 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
policyGroupId	The desired Open ID Connect policy.
grantAccessSessionRevocationApi	If set to <code>true</code> , this client is allowed to access the Session Revocation API (see Back-Channel Session Revocation on page 273).
pingAccessLogoutCapable	If set to <code>true</code> , PingFederate sends (via the browser) logout requests to an OpenID Connect endpoint in PingAccess as part of the logout process (see Asynchronous Front-Channel Logout on page 273).
logoutUris	A list of client logout URI's which will be invoked when a user logs out through one of PingFederate's SLO endpoints (see Asynchronous Front-Channel Logout on page 273). Similar to <code>redirectedUris</code> , multiple entries are allowed.

Example JSON

```
{ "client": [
  {
    "clientId": "12345",
    "name": "Client Doe",
    "refreshRolling": "true",
    "redirectUris": [
      "http://www.url.com",
      "http://www.url2.com"
    ]
    "logoUrl": "http://www.url.com/image.gif",
    "clientCertIssuerDn": "CN=CA, dc=pingidentity, dc=com",
    "clientCertSubjectDn": "cn=MyClient, dc=pingidentity, dc=com",
    "description": "Description of the client",
    "persistentGrantExpirationType": "OVERRIDE_SERVER_DEFAULT",
    "persistentGrantExpirationTime": "3",
```

```

    "persistentGrantExpirationTimeUnit": "d",
    "bypassApprovalPage": "true",
    "restrictScopes": "true",
    "restrictedScopes": [
      "scope 1",
      "scope 2"
    ]
    "grantTypes": [
      "password",
      "refresh_token"
    ]
  }
]
}

```

Returns

- 200 – Success
- 400 – Failed To Create Client
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

Description

Updates client details for a specified client.



Note: You cannot update a client ID—you must delete the client record and create a new one.

JSON Parameters

The same parameters described for *POST* on page 483 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.



Note: If the `secret` parameter is used without `forceSecretChange`, the `secret` value is ignored.

Example JSON

```

{"client": [
  {
    "clientId": "12345",
    "name": "Client Doe",
    "refreshRolling": "true",
    "redirectUris": [
      "http://www.url.com",
      "http://www.url2.com"
    ]
    "logoUrl": "http://www.url.com/image.gif",
    "forceSecretChange": "true",
    "secret": "mysecretphrase",
    "description": "Description of the client",
    "persistentGrantExpirationType": "OVERRIDE_SERVER_DEFAULT",
    "persistentGrantExpirationTime": "3",
    "persistentGrantExpirationTimeUnit": "d",
    "bypassApprovalPage": "true",
    "restrictScopes": "true",
    "restrictedScopes": [
      "scope 1",

```

```

    "scope 2"
  ]
  "grantTypes": [
    "password",
    "refresh_token"
  ]
}
] }

```

Returns

- 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


Description

Retrieves details for all OAuth clients.

JSON Parameters

None.

Returns

- 200 – Success
The body contains JSON data for all clients.
 **Note:** 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.

/oauth/clients/id

This resource accepts the methods GET and DELETE.

GET

Description

Retrieves details for the specified client ID.

JSON Parameters

None.

Returns

- 200 – Success

JSON client parameters are included.



Note: The parameter `refreshRolling` is not returned if the AS global setting is set for a client (the default).

- 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

Description

Deletes records for the specified client ID.

Returns

- 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 via this endpoint in the `runtime-api.log` file. While the events themselves are not configurable, you may 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 for OAuth Access Grant Management. This service enables retrieval and revocation of access grants for a particular client or user.

The REST resources are URL path extensions of the PingFederate runtime endpoint:

```
http(s)://<pf_runtime_host>:<runtime_port>/pf-ws/rest
```


For example: `https://my_corp.com:9031/pf-ws/rest/oauth/clients`

! **Important:** Applications must authenticate to the Web Service using HTTP Basic credentials specified in a PingFederate Password Credential Validator (see [Manage password credential validators](#) on page 196). The configured Credential Validator, in turn, must be selected in the OAuth AS configuration (see [Configure authorization server settings](#) on page 243).

Endpoints: `/oauth/clients/{clientId}/grants/{grantId}` and `/oauth/users/{userKey}/grants/{grantId}`

These resources accept the methods GET which supports retrieval and DELETE which supports revocation of access grants.

☰ **Note:** The Clients endpoint requires client records to be managed in a database, rather than in XML files (the installed default storage method). For information on setting up database management for OAuth clients, see [Define an OAuth client data store](#) on page 171.

Parameters

clientId	The client identifier to retrieve or revoke grants for (see Configure an OAuth client on page 269).
userKey	The user key to retrieve or revoke grants for. The user key is defined under IdP Adapter Mapping and Resource Owner Credentials Mapping.
grantId (optional)	Access grant identifier used to retrieve or revoke a specific grant. The value corresponds to the id field of the JSON array of grants returned from a previous GET <code>/oauth/clients/{clientId}/grants</code> or GET <code>/oauth/users/{userKey}/grants</code> . 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 request and JSON response

Request to retrieve all grants for client `im_client`:

```
GET /pf-ws/rest/oauth/clients/im_client/grants HTTP/1.1
Host: localhost:9031
Authorization: Basic YWRtaW46MkZlZGVyYXRl
X-XSRF-HEADER: PingFederate
```

Response could be:

```
{
  "items": [{
    "id": "gn3T3qGVjoHL9p8HKtCSZNriwg9H3DA8",
    "userKey": "joe",
    "grantType": "IMPLICIT",
    "scopes": [],
    "clientId": "im_client",
    "issued": "2014-03-26T21:15:38.551Z",
    "updated": "2014-03-26T21:15:38.551Z"
  }]
}
```

Returns

- 200 – Success
- 204 – Success with no content returned

Will be seen when revoking an access grant

- 401 – Invalid Credentials.
The user does not exist or the password is incorrect.
- 404 – Not found
Resource (user, client, or access grant) not found
- 500 – Internal Server Error
An unknown error has occurred.

Logging

PingFederate records the actions performed via this endpoint in the `runtime-api.log` file. While the events themselves are not configurable, you may 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 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. Note that this endpoint is not part of the OAuth specification. The **Grant Access to Session Revocation API** check box must also be selected in the configuration for the applicable clients.

 **Important:** OAuth clients must authenticate to the Web Service using their client secrets via HTTP Basic authentication or client certificates.

Endpoint: /pf-ws/rest/sessionMgmt/revokedSris

 **Tip:** Information about the Session Revocation API endpoint is also available in the OpenID Provider Configuration endpoint. Look for `ping_revoked_sris_endpoint` in the metadata.

This resource accepts the HTTP POST and GET methods. It also requires the `X-XSRF-HEADER` HTTP header with any value to prevent cross site request forgery.

 **Note:** The POST method described in this section requires the element name/value pair formatted in JavaScript Object Notation (JSON).

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 POST request and responses

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
YWNfb2ljX2NsaWVudDphYmMxMjNERUZnaGlqa2xtbm9wNDU2N3JzdHV2d3h5elpZWVdVVDg5MTBTU1FQT25tbG
X-XSRF-HEADER: PingFederate
Content-Type: application/json

{"id": "abc123"}
```

Possible HTTP response status 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 GET requests and responses

A GET request to query the revocation status for a session with a session identifier of abc123:

```
GET /pf-ws/rest/sessionMgmt/revokedSris/abc123
Host: localhost:9031
Authorization: Basic
YWNfb2ljX2NsaWVudDphYmMxMjNERUZnaGlqa2xtbm9wNDU2N3JzdHV2d3h5elpZWVdVVDg5MTBTU1FQT25tbG
X-XSRF-HEADER: PingFederate
```

If PingFederate has been configured to manage authentication sessions in the **IdP Configuration (or SP Configuration) > Sessions** screen, 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.

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; for example:

```
GET /pf-ws/rest/sessionMgmt/revokedSris/abc123&updateActivityTime=false
Host: localhost:9031
Authorization: Basic
YWNfb2ljX2NsaWVudDphYmMxMjNERUZnaGlqa2xtbm9wNDU2N3JzdHV2d3h5elpZWVdVVDg5MTBTU1FQT25tbG
X-XSRF-HEADER: PingFederate
```

Possible HTTP response status 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. As previously stated, if PingFederate has been 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 via this endpoint in the `runtime-api.log` file. While the events themselves are not configurable, you may 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 application programming interface (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 // new
- Cluster management
- Data stores and password credential validators
- Keys and certificates
- OAuth settings
- Server settings

For a complete list, see [Using the API interactive documentation](#) on page 496. For known limitations, see [Release notes](#).

Authentication

The administrative API supports various authentication options, see [Configure access to the administrative API](#) for more information.

Concurrent Access

The administrative API supports concurrent access. When concurrent API calls are made to modify the same API resource (such as the same IdP Adapter instance or the same SP connection), the last request processed by PingFederate wins.

Logging

PingFederate records actions performed via 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 is protected by:

- [Native authentication](#) — Uses credentials managed within PingFederate.
- Alternative authentication — Utilize credentials external to PingFederate through an LDAP user-data store, the RADIUS protocol, or client certificates:
 - [LDAP authentication](#)
 - [RADIUS authentication](#)
 - [Certificate-based Authentication](#)

For new installations, native authentication is chosen by default.

For upgrades, if the authentication method of the administrative API was not previously set (for example, when upgrading from PingFederate 7.3 or older), the upgrade utility sets the value to that of the administrative console; otherwise, it preserves the previously set value (for example, when upgrading from PingFederate 8.0 to a future release).

The authentication method for the administrative API can be changed at a later time to any of the four choices, regardless of which authentication method is chosen for the administrative console.

Besides authentication, 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	Admin	Configure partner connections and most system settings (except the management of native accounts and the handling of local keys and certificates).
Admin	Crypto Admin	Manage local keys and certificates.
Admin	User Admin	Create users, deactivate users, change or reset passwords, and install replacement license keys.
Auditor	<i>Not applicable</i>	View-only permissions for all administrative functions. When the Auditor role is assigned, no other administrative roles may be set.



Note: All three administrative roles are required to access and make changes through the following services:


- The `/ConfigArchive` administrative API endpoint
- The **Server Configuration** > **Configuration Archive** screen in the administrative console
- The `/pf-mgmt-ws/ws/ConnectionMigrationMgr` and `/pf-mgmt-ws/ws/ConfigReplication` Connection Management Web Service endpoint
- The **Server Configuration** > **Application Authentication** screen in the administrative console for the **Connection Management** service

Enable native authentication


When the administrative API is protected by native authentication, access to the administrative API is restricted to the users defined in the **Account Management** screen. 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.

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. Log on to the administrative console with an account that has the role `User Admin`.

 **Important:** When the administrative console is protected by an alternative console authentication (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 **Account Management** functionality unless the logged-on administrator has been granted the `User Admin` right.

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 such administrator logs on to the administrative console, the **Account Management** screen becomes available for her or him 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 112.

3. Click **Server Configuration** on the Main Menu.
4. Click **Account Management** under Administrative Functions.
5. Create or manage users as needed, including assigning various PingFederate administrative roles as indicated by the **PingFederate User Access Control** table in [Configure access to the administrative API](#) on page 493.

 **Note:** When assigning role(s), keep in mind that all users defined in the **Account Management** screen can be used to access the administrative API and the administrative console.

Enable 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. The LDAP authentication setup, including role assignment, is available via `<pf_install>/pingfederate/bin/ldap.properties`. The roles assigned to the LDAP accounts affect the results of the API calls.


 **Note:** When LDAP authentication is configured, PingFederate does not lock out accounts based upon the number of failed logon attempts. Responsibility for preventing access is instead delegated to the LDAP server and enforced according to its password lockout settings.


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.


See the comments in the file for instructions and additional information.

 **Important:** Be sure 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 493.

 **Note:** When assigning role(s), keep in mind that all LDAP accounts specified in `ldap.properties` can be used to access the administrative API and the administrative console.

 **Tip:** You can also use this configuration file in conjunction with RADIUS authentication to determine permissions dynamically via an LDAP connection.


3. Restart PingFederate.


 **Note:** In a clustered PingFederate environment, you only need to modify `run.properties` and `ldap.properties` on the console node.

Enable RADIUS authentication

The RADIUS authentication setup is available via configuration files in the `<pf_install>/pingfederate/bin` directory. The RADIUS protocol provides a common approach for implementing strong authentication in a client-server configuration. The administrative API supports the protocol scenario for one-step authentication (for example, appending an OTP after the password).

When the administrative API is protected by RADIUS authentication, 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 RADIUS authentication is configured, PingFederate does not lock out accounts based upon the number of failed logon attempts. Responsibility for preventing access is instead 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.


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.

See the comments in the file for instructions and additional information.

 **Important:** Be sure to 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 (also in the `bin` directory) to map LDAP group-based permissions to PingFederate roles. (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 493.)

 **Note:** When assigning role(s), keep in mind that all accounts specified in `radius.properties` can be used to access the administrative API and the administrative console.

3. Restart PingFederate.

 **Note:** In a clustered PingFederate environment, you only need to modify `run.properties` and `radius.properties` on the console node.


Enable 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. In addition, the corresponding root CA certificate(s) must either be contained in the Java runtime or be imported into the PingFederate's Trusted CA store (see *Manage trusted certificate authorities* on page 178).

The rest of the certificate-based authentication setup, including specifying the Issuer DN of the root CA certificate(s) and the applicable role(s) of the client certificate(s), is available via `<pf_install>/pingfederate/bin/cert_auth.properties`. The roles assigned to the certificates affect the results of the API calls.

1. Log 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).

To import a certificate, click **Trusted CAs** in the Certificate Management section under Server Configuration.


 **Tip:** You may wish to click the Serial number and copy the Issuer 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 (see the properties file for more information).

 **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 role(s), 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 493.

 **Important:** The configuration values are case-sensitive.


 **Note:** When assigning role(s), 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.

Using 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. In general, the API calls can be called from an interactive user interface, custom applications, or from command line tools such as `cURL`. The endpoint is only available on the `pf.admin.https.port` defined 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/POST requests.

To access the interactive documentation in PingFederate

1. Start PingFederate.
2. Launch your browser and go to:


`https://<DNS_NAME>:9999/pf-admin-api/api-docs/`

where `<DNS_NAME>` is the fully qualified name of the machine running the PingFederate server.

To test an administrative API using the interactive documentation:


1. Click a section of the administrative API you would like to explore—for example, `/sp/idpConnections`.
2. Expand the method you want to use—for example, GET.
3. Enter any required parameters.
4. Click **Try it out!**

The Request URL, Response Body, Response Code, and Response Headers appear.


 **Note:** You may be prompted to log in using administrative credentials over HTTP Basic Authentication. Enter the credentials of any existing administrator. The role of the administrator affects their access to the requested API.

Attribute mapping expressions

PingFederate® provides an advanced option allowing administrators to map user attributes by way of an expression language. Because the option carries with it a potential for misuse, however, it is disabled in the administrative console for security reasons.


 **Tip:** If you are upgrading to PingFederate 5.1 or higher and importing a configuration archive that uses expression mapping, the feature is enabled automatically.

This topic describes the option, which is based on the Object-Graph Navigation Language (OGNL), and how to enable or disable it.


 **Caution:** The security concern posed by expressions is related to a potential for abuse by PingFederate administrative users within an organization; the concern is not related to any known external threats. We recommend, however, that the option be enabled only if required.

Enable and disable expressions

An administrator can manually enable or disable the use of expressions in PingFederate® by editing a configuration file.


 **Important:** If the current configuration contains expressions, disabling the feature causes errors during runtime processing.

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; for example:

```
<?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 a value (the installed default) *does not* necessarily disable the use of expressions. To facilitate backward compatibility, when no value is present, configuration archives containing expressions can be imported successfully, and further use of the feature is enabled. (The term “silent” is used for this condition in the server log.)

3. Start or restart PingFederate.

 **Tip:** If you are enabling expressions to use for mapping outbound provisioning attributes, it is not necessary to restart the PingFederate server.

4. If you have a clustered PingFederate environment:

- a) Sign on to the PingFederate administrative console.
- b) On the **Server Configuration > Cluster Management** screen, click **Replicate Configuration**.
- c) Restart PingFederate on the engine nodes.

When you enable expressions, these expressions are available for use in multiple locations:

- The **Source** list under each of the administrative-console contract fulfillment screens.
- The **Show Advanced Criteria** section on the **Issuance Criteria** screen following each of the administrative-console contract fulfillment screens.
- The provisioning attribute-mapping screen (when the **Outbound Provisioning** role is enabled).

When you make this selection, you can enter the expression in the text field provided. You can also test expressions (see [Using the OGNL edit screen](#) on page 501).

Construct OGNL expressions

OGNL is based on the Java programming language. OGNL expressions are useful for evaluating and manipulating attribute values and returning information based on the results. You can also 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 IdP, PingFederate provides predefined OGNL variables for IdP-adaptor attributes, any attributes retrieved from data stores, and attributes for token authorization. For an SP, variables are available for attributes received in an assertion, an attribute query, and attributes for token authorization. For example, the SAML_SUBJECT value may be retrieved using:

```
#SAML_SUBJECT
```



Note: Use the following construction for any attributes from any source that contain special characters (hyphens, for example), which cannot be parsed by OGNL: `#this.get("<attribute_name>")`



Note: 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:

```
#SAML_SUBJECT="usr@msn.com"
not:
#SAML_SUBJECT=usr@msn.com
```

For more information, see [Using the OGNL edit screen](#) on page 501.

For more information about OGNL, including detailed user documentation, see the Apache Commons OGNL page (commons.apache.org/ognl).

Data store syntax

For data-store attributes with an attribute source ID, use this syntax:

```
#this.get("ds.attr-source-id.attribute_name")
```

For data-store attributes without an attribute source ID, use this syntax:

```
#this.get("ds.attribute_name")
```

Issuance criteria syntax

To access mapped attributes when configuring token-authorization expressions, use this syntax:

```
#this.get("mapped.attribute_name")
```

To access most context attributes when configuring token-authorization expressions, use this syntax:

```
#this.get("context.attribute_name")
```

To access the HTTP Request context attribute, use this syntax:

```
#this.get("context.HttpServletRequest").getObjectValue()
```

The returned value will be an instance of `javax.servlet.http.HttpServletRequest` (see docs.oracle.com/javaee/7/api/javax/servlet/http/HttpServletRequest.html).

Sample OGNL expressions

Below are examples of using OGNL expressions for attribute mapping and token authorization.

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"
```

Multi-valued attribute

```
new org.sourceid.saml20.adapter.attribute.AttributeValue( {"Blue", "Gray",
"Pink"})
```

This expression formulates a multi-valued attribute in an SSO token; for example:

```
<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 multi-valued attribute. The former is a SAML assertion via a SAML SP connection. The latter is a JSON Web Token (JWT) via a WS-Federation SP connection using JWT as the token type.

Token authorization

This expression verifies whether a user is a member of the “Contractor” or “Employee” group.

```
#this.get("ds.ldap.memberOf") != null?
(
  #this.get("ds.ldap.memberOf").toString().matches("(?i)
  .*CN=Contractor,cn=users,dc=company,dc=com.*|
  .*CN=Employee,cn=users,dc=company,dc=com.*")
):@java.lang.Boolean@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; statements calling methods whose arguments are enclosed in quotes must be entered 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::74
```

The `isInRange` method supports both IPv4 and IPv6 CIDR notations.

HTTP request context

The following example may be used to retrieve a value from an HTTP request object. In this case, the expression retrieves the `User-Agent` HTTP header value and compare it against a value required for token authorization.

```
#this.get("context.HttpRequest").getObjectValue().getHeader("User-
Agent").equals("somevalue")
```

STS client authentication context

This STS SSL Client Certificate Chain example checks that the issuer of the client certificate matches the specified 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 that `#this.get("context.StsSSLClientCertChain").getObjectValue()` returns an array of `java.security.cert.X509Certificate` instances (see docs.oracle.com/javase/8/docs/api/java/security/cert/X509Certificate.html); this array starts with the client certificate itself.

Issuance criteria and multiple virtual server IDs

When you use virtual server IDs to connect to multiple environments in one connection (see [Connecting to a partner in multiple connections](#) on page 84), 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 IdP) have their own departmental subdomains, `sales.contoso.com` and `support.contoso.com`. The SP identifies both environments under the parent domain, `contoso.com`.


In this scenario, the PingFederate IdP server can be configured 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. For example:

```
#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; statements calling methods whose arguments are enclosed in quotes must be entered on a single line.


Using the OGNL edit screen

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.

 **Note:** For information about using OGNL, please refer to the commons.apache.org/proper/commons-ognl/language-guide.html.

- To reach the OGNL editor, click **Edit** under **Actions** for an expression on any of the **Attribute Fulfillment** screens or click **Test** in the **Show Advanced Criteria** section on the **Issuance Criteria** screen.


The **Edit** action is also available on Issuance Criteria screens for expressions (see [About token authorization](#) on page 74).

 **Note:** The test function does not work for the `context.httpRequest` attribute, because its value is an object rather than text.

Here is an example of the edit screen, from the IdP configuration flow:

- 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 screen.

If the expression contains no errors, the result appears under **Test Results**.

 **Important:** If you make changes to an expression and want to save it, click **Update** under **Actions**. To discard changes, click the **Cancel link** under **Actions**; clicking the **Cancel** button near the bottom of the screen discards all changes you made in the current task.

Customize assertions and authentication requests


Some Browser SSO use cases may require additional customizations in the assertions sent from the PingFederate® IdP server to the SP or the authentication requests sent from the PingFederate SP server to the IdP. PingFederate is capable of fulfilling these use cases on a per-connection basis using OGNL expressions.

1. If you have already done so, enable OGNL expression by editing the `org.sourceid.common.ExpressionManager.xml` file in the `<pf_install>/pingfederate/server/default/data/config-store` directory.
2. Select the applicable SP or IdP connection.
3. On the **Activation & Summary** screen, scroll down to the **Protocol Settings** section, and select the **URL** screen.
4. Click **Show Advanced Customizations** to begin customizing the applicable message.

The available **Message Types** that can be customized varies depending your federation role (IdP or SP) as well as the protocol of the connection (SAML 1.x, SAML 2.0, and WS-Federation). Once a message type is selected, you have access to a list of the **Available Variables**. By calling various methods, you can customize the assertions or the authentication requests to fulfill your use case.

Message types and available variables

The following tables describe the relationship between message type and available variable, as well as the corresponding class or interface information in Javadoc.

 **Tip:** The Javadoc for PingFederate is located 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.RequestSecurityTokenResponseDocument

Message Types	Available Variables
	Classes/Interfaces in Javadoc #Attributes org.sourceid.util.log.AttributeMap

IdP connections (SAML 2.0)

Message Type	Available Variables
	Classes/Interfaces in Javadoc #AuthnRequestDocument org.sourceid.saml20.xmlbinding.protocol.AuthnRequestDocument

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
#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.sourceid.saml20.xmlbinding.protocol.AuthnRequestDocument
SP connection	SAML 2.0	#FedHubOutgoingAuthnRequest org.sourceid.saml20.xmlbinding.protocol.AuthnRequestDocument
SP connection	SAML 2.0	#FedHubIncomingAuthnResponse
	SAML 1.x	org.sourceid.saml20.xmlbinding.protocol.ResponseDocument (SAML 20)
	WS-Federation	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.

Connections	Protocol	Available Variables Classes/Interfaces in Javadoc
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

Below are examples of using OGNL expressions to customize assertions and authentication requests.

Add SessionNotOnOrAfter to assertions

This expression adds the optional SessionNotOnOrAfter attribute in the <AuthnStatement> element and set 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:

```
...
<saml:AuthnStatement ... AuthnInstant="2015-03-20T16:27:37.344Z"
  SessionNotOnOrAfter="2015-03-20T17:27:37.398Z">
  <saml:AuthnContext>
    <saml:AuthnContextClassRef>...</saml:AuthnContextClassRef>
  </saml:AuthnContext>
</saml:AuthnStatement>
...
```

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.

In this example, we use 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** screen. You are free to use other application logic in this workflow.

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
  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
    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:Attribute Name="ExtAttr1"
  NameFormat="urn:oasis:names:tc:SAML:2.0:attrname-format:unspecified">
  <saml:AttributeValue xsi:type="xs:string"
    xmlns:xs="http://www.w3.org/2001/XMLSchema"
    xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance">
    123
  </saml:AttributeValue>
</saml:Attribute>
<saml:Attribute Name="ExtAttr2"
  NameFormat="urn:oasis:names:tc:SAML:2.0:attrname-format:unspecified">
  <saml:AttributeValue
    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>
...

```

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 SP-initiated SSO request.

Message Type: `AuthnRequestDocument`

Expression:

```

#element = #XmlHelper.addToSaml2Extensions(#AuthnRequestDocument,
  '<samlens:orgId name="orgId" xmlns:samlens="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://<pf_host>/sp/startSSO.ping?PartnerIdpId=<Idp>&oid=123` would trigger the following `Extensions` block:

```

<samlp:AuthnRequest ...>
  <saml:Issuer ...>...</saml:Issuer>
  <samlp:Extensions>
    <samlens:orgId name="orgId" value="123"
  xmlns:samlens="urn:org.sample.wms"/>
  </samlp:Extensions>
  ...
</samlp:AuthnRequest>

```

Fulfillment by data store queries


PingFederate® can be configured to use data stores to supply attributes for various contracts (for example, the IdP's adapter contract, IdP's attribute contract, the SP's adapter contract, or STS token contracts) or conditions to be verified in the token authorization process. Standard data stores may include any JDBC-accessible database or an LDAP v3-compliant directory server. Alternatively, you can create your own driver for non-JDBC or non-LDAP data stores, such as flat files or SOAP-connected databases, using the PingFederate Custom Source SDK, and then map from custom fields to fulfill your use cases.

Attribute mapping with multiple data sources

When querying data stores for attributes to help fulfill a contract on the IdP side, PingFederate® can be configured to use more than one attribute source.

Multiple data stores in one mapping

The PingFederate IdP server supports separate data stores to look up attributes for a single mapping. For example, you can query multiple data stores for values and map those values in one mapping, or query a data store for a value and use that value as input for subsequent queries of other data stores.



If a data store 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 data store 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) cannot be fulfilled, the request fails.

Multiple data stores in one mapping is available for Browser SSO and WS-Trust STS on the IdP side, authentication policy contract to SP adapter mappings, and the following OAuth workflows:

- OpenID Connect policies (the user-attributes mapping process)
- Resource-owner credential mappings
- IdP adapter mappings
- Access token mappings

Multiple mappings and failsafe mapping

For Browser SSO and WS-Trust STS on the IdP side, PingFederate also supports separate search parameters for each data store and for "fall-through" searches. For example, you can add the same data store more than once, using different search queries for each instance, or you can search different data stores successively. If any search fails to find a user in the specified attribute source, the next search is executed until a match is found. A failsafe attribute source can also be configured, providing a default set of attributes from the IdP adapter and using text values.



Data store query configuration

To query attributes from a data store, you must

- Choose a data store instance
- Specify the search location and the desired attributes
- Enter the search term to find the end users

Once the configuration is complete, you can fulfill a contract or verify a condition in the token authorization workflow using the results from the data store queries.

Choose a data store

On the **Data Store** screen, choose a data store for PingFederate® to look up attributes.

1. Enter a description (and ID if prompted) for the data store.
2. Select a data store instance from the **Active Data Store** list.

i **Tip:** 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.

3. Depending on the data store type, the rest of the setup varies as follows:


Data store type	Required tasks
JDBC	<ul style="list-style-type: none"> • <i>Specify a JDBC database table and columns</i> on page 508 • <i>Enter a database search filter</i> on page 509
LDAP	<ul style="list-style-type: none"> • <i>Specify an LDAP directory and attributes</i> on page 510 • <i>Define encoding for LDAP binary attributes</i> on page 511 (optional) • <i>Enter an LDAP search filter</i> on page 511
Custom	<ul style="list-style-type: none"> • <i>Specify a custom source filter and fields</i> on page 512

Specify a JDBC database table and columns

When you choose to use a database source for attributes, you follow this path through the configuration steps.

On the **Database Table and Columns** screen you begin to specify exactly where additional data can be found to fulfill your use case. Only one table may be used as a source of data for a JDBC query. For more information about each field, refer to the following table:

Field	Description
Schema	Lists the table structure that stores information within a database. Some databases, such as Oracle, require selection of a specific schema for JDBC queries. Other databases, such as MySQL, do not require selection of a schema.
Table	Displays the tables contained in the database. Select the table to retrieve data from the data store.
Columns to return from SELECT	Displays selected columns from the selected tables. Select the columns that are associated with the desired attributes you would like to return from the JDBC queries.


 **Important: (For MySQL users)** To allow for table and column names that may 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 connection string of your JDBC data store instance; for 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 dev.mysql.com/doc/refman/5.7/en/server-system-variables.html.

1. Choose a schema (when applicable) from the **Schema** list.
2. Select a table from the **Table** list.
3. Optional: Click **View Attribute Contract** to determine what attributes to look up.
4. Optional: Click **Refresh** if you are updating an existing configuration and changes may have been made to the database.
5. Under **Columns to return from SELECT**, choose a column name and click **Add Attribute**.

Repeat as needed to add more columns.

 **Note:** It is not required to add a column here for it to be used as part of a search filter. Add only the attributes required for contract fulfillment or token authorization.

Example

Suppose you (the IdP) have a data table named **ACCESSTABLE** with three columns: `userid`, `department`, and `accesslevel`. Your use case requires `accesslevel` to be mapped into a SAML contract to an SP.

On the **Database Table and Columns** screen, select the **ACCESSTABLE** table and add the `accesslevel` column.

Enter a database search filter

On the **Database Filter** screen, enter a JDBC `WHERE` clause for PingFederate® to query the database table you selected to retrieve a record associated with a particular value (or 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** screen.

 **Tip:** To get a list of columns, click the **View List of Columns from ...** link.

The right side (`value1`) is the match-against value, generally a variable passed in from either an authentication source (for an IdP) or an assertion (for an SP). The variables are shown underneath the **Where** 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 screen.

You can also apply additional search criteria from your own database, using any other columns from the targeted table.

1. Enter a *WHERE* clause in the text field. The initial *WHERE* is optional.
2. Ensure the syntax and variable names are correct. For more information about *WHERE* clauses, consult your DBMS documentation.
3. Click **Next** to complete the configuration to query attributes from a JDBC data store.

Later in the workflow, you can use the attribute values returned from the database in the applicable contract fulfillment screen, the issuance criteria screen, or both, to fulfill your use case.

Example

Suppose you have selected a data table named **ACCESSTABLE** on the **Database Table and Columns** screen. 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** screen, 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.

Specify an LDAP directory and attributes

When you choose to use an LDAP source for attributes, you follow this path through the configuration steps.

On the **LDAP Directory Search** screen you begin to specify the branch of your LDAP 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 of the tree structure in which the search begins. This field is optional if records are located at the LDAP root.
Search Scope	The node depth of the query. Select Subtree (the default value), One level or Object.
Root Object Class	The class containing the desired attributes.
Attributes	A list of attributes based on the selected Root Object Class value. Subject DN, which may be used as the primary user identifier, is always returned by default.

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, select an attribute, and then click **Add Attribute**.

It is not required to add an attribute here for it to be used as part of a search filter. Add only the attributes required for contract fulfillment or token authorization.

Repeat this step to add more attributes as needed.



Note: When connecting to 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, suppose 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.)

For Oracle Directory Server, choose isMemberOf under **Attribute** for nested group membership. For more information, see [documentation about isMemberOf from Oracle](https://docs.oracle.com/cd/E29127_01/doc.111170/e28967/ismemberof-5dsat.htm) (docs.oracle.com/cd/E29127_01/doc.111170/e28967/ismemberof-5dsat.htm)



Tip: 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.

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 LDAP server, OU=West, DC=example, DC=com.

On the **LDAP Directory Search** screen, 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**.

Define encoding for LDAP binary attributes

The **LDAP Binary Attribute Encoding Types** screen appears when a binary attribute is added on the **LDAP Directory Search** screen. (Attributes are specified as binary in the **Server Configuration > Data Stores > the applicable LDAP connection > LDAP Configuration > Advanced** screen.) Since binary attribute data cannot be used in an assertion to the SP, specify the encoding type (**Base64**, **Hex**, or **SID**) that you want to apply during fulfillment.



Note: Defining encoding for LDAP binary attributes is only applicable to IdP and IdP-to-SP bridging use cases.

To set an encoding type, select a value from the **Attribute Encoding Type** list.

Repeat this step for each LDAP attribute that has been specified as binary in the LDAP connection.

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.

Enter an LDAP search filter

On the **LDAP Filter** screen, enter a LDAP filter for PingFederate® to query the data you selected to retrieve a record associated with a particular value (or values) from the user's session. The filter is in the form:

```
attribute1=value1
```

The left side (*attribute1*) is an attribute from your LDAP directory.

 **Tip:** To get 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 IdP) or an assertion (for an 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 screen.

You can also apply additional search criteria from your own LDAP server, using any other attributes from the target object class.

In general, a filter narrows a search to locate requested data by either including or excluding specific records. An LDAP 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) to search on, a search filter operator that will determine what to match, and the value of the attribute being sought.

1. Enter an LDAP search filter in the text field.
2. Ensure the syntax and variable names are correct. For general information about search filters, consult your LDAP documentation.
3. Click **Next** to complete the configuration to query attributes from an LDAP data store.

Later in the workflow, you can use the attribute values returned from your LDAP server in the applicable contract fulfillment screen, the issuance criteria screen, or both, to fulfill your use case.

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 LDAP 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** screen, 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.

Specify a custom source filter and fields

When you choose to use a custom source for attributes, you follow this path through the configuration steps.

On the **Configure Custom Source Filter** screen you begin to specify a filter, or search query, for your custom data source. This screen display and the syntax of the filter depends on your developer's implementation of the custom source SDK.


On the **Configure Custom Source Fields** screen, 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.


Later in the workflow, you can use the values returned from the custom source in the applicable contract fulfillment screen, the issuance criteria screen, or both.

Configure failsafe options

When a data store is configured and the attribute mappings under **Attribute Sources & User Lookup** fail to complete the attribute contract in an 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** screen does not appear if you chose to retrieve attributes from multiple data stores using a single mapping on the **Mapping Method** screen.

 **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 via the data-store mapping setup but others are not.

The failsafe mapping is used only when *all* of the mappings configured in the data-store 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.


- 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** screen.
- To stop the SSO transaction, select **Abort the SSO transaction**.

Review data store query configuration

You can review the configuration on the **Summary** screen.

To change your configuration, click the heading over the information you want to edit.

When you are finished, click **Done** to continue with the rest of the configuration, where you can use the attribute values returned from the data store in the applicable contract fulfillment screen, the issuance criteria screen, or both, to fulfill your use case.

 **Note:** Remember to click **Save** at the end of your workflow to save your settings.


Troubleshooting

Basic troubleshooting tips are provided here to help overcome common difficulties. Help is also available from the [Support Center](#).

Enable debug messages and console logging

Starting with version 8.2, PingFederate only records messages that are tagged with log level `INFO`, `WARN`, `ERROR`, and `FATAL` to the server log (and the provisioner log). Messages that are tagged `DEBUG` (or `TRACE`) are not recorded to optimize performance. Console logging is also disabled for the same reason.

For troubleshooting purpose, you may adjust the log level to `DEBUG` in the `log4j2.xml` file and optionally re-enable console logging.

 **Important:** When debug messages and console logging are no longer required, ensure they are turned off.

1. Edit the `<pf_install>/pingfederate/server/default/conf/log4j2.xml` file.
 - a) To enable debug messages, look for the `Limit categories` section, and then update from

```
<Logger name="org.sourceid" level="INFO" />
```

to

```
<Logger name="org.sourceid" level="DEBUG" />
```

b) To enable console logging, look for the Set up the Root logger section, and then update from

```
<AsyncRoot level="INFO" includeLocation="false">
  <!-- <AppenderRef ref="CONSOLE" /> -->
  <AppenderRef ref="FILE" />
</AsyncRoot>
```

to

```
<AsyncRoot level="INFO" includeLocation="false">
  <AppenderRef ref="CONSOLE" />
  <AppenderRef ref="FILE" />
</AsyncRoot>
```

2. Save any changes made.

For a clustered PingFederate environment, repeat these steps on each applicable node.

Updating the log level for existing `Logger` elements does not require a restart. The adjustments are activated within half a minute. Other changes require a restart of PingFederate; for example:

- Adding new logger elements
- Updating the `size` attribute in the `RollingFile/Policies/SizeBasedTriggeringPolicy` element
- Updating the `max` attribute in the `RollingFile/DefaultRolloverStrategy` element
- Updating the `filePattern` attribute without modifying the `fileName` attribute in the `RollingFile` appender

Server startup

Problem	Solution
PingFederate does not start.	Make sure that Oracle Java SE Runtime Environment (Server JRE) is installed (see Install Java on page 19 in the “Installation” chapter of the <i>Get started with PingFederate</i> guide).
The server starts but indicates the license file is not found or invalid.	Ensure a current license is installed (see Manage PingFederate license on page 91).

Data store issues

Problem	Solution
When setting up the JDBC data store, a connection cannot be established.	<p>Verify that the proper drivers and connectors have been installed.</p> <p>Also, verify the connection URL, username, and password. If unsuccessful, contact your database administrator.</p> <p>For more information, see Configure a JDBC database connection on page 159.</p>
Cannot connect to a Directory Service with the LDAP protocol.	<p>Verify the connection URL, port, principal, and credentials. If unsuccessful, contact your system administrator (see Configure an LDAP connection on page 161).</p> <p>If using LDAPS, ensure the LDAP server's SSL certificate is signed by a trusted certificate authority or is imported into PingFederate (see Manage trusted certificate authorities on page 178).</p>

Resolve URL-related errors

If a user encounters a `404 Not Found` status or a `System Error` message, check the URL of the request.


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 (?).

 **Tip:** 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](http://www.w3schools.com/tags/ref_urlencode.asp) (www.w3schools.com/tags/ref_urlencode.asp).

The remedy for both sample issues is to update the URL of the IdP-initiated SSO to:

`https://sso.idp.local/idp/startSSO.ping?PartnerSpId=sp1&TargetResource=https%3A%2F%2Fapp.sp1.local%2F`

Resolve service-related errors

If a user encounters an `Unexpected System Error` message with a reference code, ask the user for the reference code and search the value in the server log. The log message should help determining the root cause, which usually requires a configuration change.


If a user encounters a `partner not active` status, select `Active` in the **Connection Status** field and click **Save** in the **Activation & Summary** screen for the connection.

Example

Unexpected System Error

When a PingFederate® IdP server receives a SAML AuthnRequest message via the Redirect binding but such SAML profile is not selected in the applicable 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
```

 **Tip:** In this sample log message, `kwlqbn` is the reference code.

The remedy is to update the applicable SP connection to allow the Redirect binding for inbound messages from the SP if the Redirect binding is one of the mutually-agreed SAML bindings that both parties should use. Alternatively, the SP can send SAML AuthnRequest messages via an allowable SAML binding based on the configuration of your SP connection.

Troubleshoot runtime errors

Users may encounter runtime errors when trying to connect to your partners. To troubleshoot these errors, investigating the user activities (in terms of the requests and the responses between the client and the PingFederate® server) helps determining 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* (www.telerik.com/fiddler) and *Charles* (www.charlesproxy.com).

Server side

Review server log messages to investigate what PingFederate has received from the client.

On rare occasions, you may also use third-party tools, such as *Wireshark* (www.wireshark.org), or tools from the operating systems, such as *tcpdump* (www.tcpdump.org), to capture network traffic on the PingFederate server.

 **Tip:** For instructions, please refer to the documentation of the third-party tools.


To correlate server log messages to user activities, you can use one of the following factors:

- The tracking ID, which can be configured to be shown in the user-facing error pages for PingFederate 7.2 (or higher).
- 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.

Activate tracking ID in templates

You can configure PingFederate® 7.2 (or higher) to display the tracking ID in the user-facing error Velocity templates. When an error occurs, you can use the tracking ID to look for the related log messages. 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`

 **Note:** You can find these Velocity template files in the `<pf_install>/pingfederate/server/default/conf/template` directory.

1. Open the applicable Velocity template file.
2. Search for the `$TrackingId` variable.
3. Follow the inline instructions to activate the variable.

Template customization does not require a restart of PingFederate. For a clustered PingFederate environment, repeat these steps on each engine node.

The following screenshot 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:V3IwuUsy8PQp-9ZbE9UfUjOEo9c



Correlate 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 locating related log messages and payloads for a given transaction for troubleshooting.

1. Ask the user for the **Tracking ID** value in the error message.
2. Search for the tracking ID in the server log, for example:
3. Use the tracking ID to review log messages and payloads pertaining to this transaction.

Generally speaking, log messages that are tagged with `WARN` or `ERROR`, or prefixed with `Caused by` are most useful.

Example

Suppose an error had 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=OaxBwPGw50BeHVXe1sgifB7iZR5Rz2VI4rhJwqUSIXV; 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 IdP adapter instance.

Correlate log messages by PF cookie

If you are using PingFederate® 7.1 (or older), or 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.

1. Capture HTTP traffic and look for the PF cookie value.

2. Search for the PF cookie value in the server log.
3. As 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.

Generally speaking, log messages that are tagged with `WARN` or `ERROR`, or prefixed with `Caused by` are most useful.

Example

Suppose an error had occurred and the associated the PF cookie value was `OaxBwPGw50BeHVXe1sgifB7iZR5Rz2VI4rhJwqUSIXV`. 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=OaxBwPGw50BeHVXe1sgifB7iZR5Rz2VI4rhJwqUSIXV; 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.
```

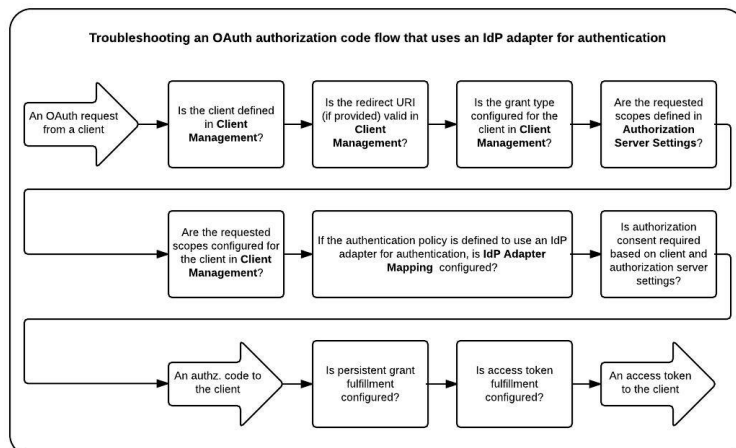
...

```
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 IdP adapter instance.

Troubleshooting OAuth transactions

Troubleshooting OAuth use cases involves reviewing the OAuth requests and various OAuth settings. This troubleshooting guide walks through an OAuth request by inspecting the parameters provided by the client in different stages of an OAuth transaction, and then compares the parameter values against the corresponding settings defined in the PingFederate® administrative console.



While the guide focuses on an OAuth authorization code use case, in which the end user authenticates through an IdP 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 may fail at which stage.

This troubleshooting guide is intended for administrators who are familiar with PingFederate and OAuth in general.

Review an OAuth request and various OAuth settings

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

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.

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 the **OAuth Settings > Client Management** screen.
- b) Look for the client by its ID.

If the client is not found, PingFederate denies the request, returns a 400 error to the client, and logs an Unknown or invalid client_id message to the server log, similar to the following:

```
2015-11-29 22:11:35,812 tid:aUBgyTMjPfuSFp5BYtsK6Cb2McM DEBUG
[org.sourceid.oauth2.handlers.HandleAuthorizationRequest]
Normal exception being handled during OAuth request processing:
org.sourceid.oauth2.handlers.AuthorizationRequestException: Unknown or
invalid client_id
```

3. Check if the redirect_uri parameter value is valid for the client.

- a) Go to the **OAuth Settings > Client Management** screen.
- b) Select the applicable client.
- c) Compare the `redirect_uri` parameter value against the values defined in the **Redirect URIs** field.

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:

```
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 the **OAuth Settings > Client Management** screen.
- b) Select the applicable client.
- c) Verify the `response_type` is selected in the **Allowed Grant Types** field.

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:

```
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 (via the `scope` parameter) are defined for the authorization server.

- a) Go to the **OAuth Settings > Authorization Server Settings** screen.
- b) Compare the scopes requested against the values defined in the **Scope Value** or the **Scope Group Value** fields.

If the scopes requested are not defined, 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, similar to the following:

```
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 the **OAuth Settings > Client Management** screen.
- 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.

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 IdP adapter for authentication. Check the IdP adapter mapping and the runtime selection made by the end user.

- a) Go to the **OAuth Settings > IdP Adapter Mapping** screen.
- b) Verify an entry exists for the IdP adapter involved.

If more than one option is available, authentication policies may be used to select an authentication source. If no authentication policy is defined or applicable, the end user is prompted with a list of all available authentication sources. The end user also has the option to save the preferred authentication source for later (in the form of a `pfidpaid` cookie).

If 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 end user an authorization consent page unless a bypass option is configured. Review the authorization approval settings.

- a) Go to the **OAuth Settings > Authorization Server Settings** screen.
- b) Review the **Reuse Existing Persistent Access Grants for Grant Types** setting.

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) Go to the **OAuth Settings > Client Management** screen.
- d) Select the applicable client.
- e) Review the **Bypass Authorization Approval** setting.

If the **Bypass Authorization Approval** check box is selected, the authorization consent page is bypassed as well.

9. 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** screen; this is the first stage of the two-stage OAuth attribute mapping process. Verify a mapping is configured.

In this example, because the end user authenticates via an IdP adapter, check the IdP adapter mapping.

- a) Go to the **OAuth Settings > IdP Adapter Mappings** screen.
- b) Verify an entry exists for the IdP adapter involved and review its configuration.

10. 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 AS uses the following rules to determine which ATM instances it should use:

1. Limit 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** screen. For OAuth Assertion Grant requests, it is the set of instances for which a mapping is defined in the IdP connection. Furthermore, the ACL (if configured) 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 RS clients, you may 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 RS clients do not provide either the `access_token_manager_id` or `aud` parameter in their requests and the RS clients have not been configured to validate against any eligible ATM instances, the same logic applies.

Review the request and the settings related to access token management.

- 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 the **OAuth Settings > Client Management** screen.
- c) Select the applicable client.
- d) Verify if a default access token is selected from the **Default Access Token Manager** list.
- e) Go to the **OAuth Settings > Access Token Mapping** screen.
- f) Review the attribute mapping configuration for the authentication source (if such mapping exists) or the `Default` mapping.

Other runtime issues

Problem	Solution
Certificates unexpectedly expire.	Verify that the server clocks are synchronized on both sides of the federation. Note that you can configure PingFederate to notify administrators in advance of impending certificate expiration (see <i>Configure runtime notifications</i> on page 146).
Receive <code>CrossModule</code> or <code>Network</code> error messages when PingFederate is deployed with a supported HSM.	Verify network connections to the Hardware Security Modules (HSMs) are active and running. Also ensure the HSMs have not been unintentionally shut down.

Glossary

access token

A data object by which a client authenticates to a Resource Server and lays claim to authorizations for accessing particular resources.

account link

A persistent name identifier that enables federation of separately established accounts among disparate domains (see also *account linking* and *pseudonym*).

account linking

A form of identity mapping among separate user accounts managed under different domains. The mapping typically involves a name identifier—which may be a pseudonym—used to link the user to each account. The identifier is persisted at the SP site to enable seamless SSO/SLO. Additional attributes may be sent with the identifier.

account mapping

A form of identity mapping by which one or more user attributes is passed in a single sign-on transaction. The attributes are used at the destination site as a means identifying the user and looking up local account information.

adapter

Plug-in software that allows PingFederate to interact with web applications and authentication systems (see *SSO integration kits and adapters* on page 67).

adapter contract

A list of attributes “hard-wired” to an adapter and conveyed generally via cookies between the adapter and application.

artifact

A reference to a SAML protocol message. The federation partner that receives the artifact dereferences it, identifying the sender, and requests the complete message in a separate SOAP transaction.

Artifact Resolution Service

The SOAP endpoint that processes artifacts returned from a federation partner to retrieve the referenced XML message. Can be used to dereference authentication requests, assertion responses, and SLO messages.

assertion

A SAML XML document that contains identifying information about a particular subject; i.e., a person, company, application, or system. A SAML assertion can contain authentication, authorization, and attribute information about the subject.

Assertion Consumer Service

A SAML-compliant portion of PingFederate in an SP role that receives and processes assertions from an IdP.

attributes

Distinct characteristics that describe a subject. If the subject is a website user, attributes may include a name, group affiliation, email address, etc.

attribute contract

A list of attributes, agreed to by the partners in an identity federation, representing information about a user (SAML subject). The attributes are sent from the IdP to the SP during SSO or STS processing.

attribute mapping

A form of identity mapping between IdP and SP user accounts that uses attributes to identify the user or provide supplemental information.

attribute source

An data source used to fulfill a requestor's attribute contract.

audience

The XML element in a SAML assertion that uniquely identifies a Service Provider.

authentication context

An element in a SAML assertion indicating the method or process used by an IdP to authenticate the subject of the assertion; may be used for authorization decisions or auditing compliance.

Authentication Request (OpenID Connect)

An OAuth 2.0 Authorization Request using extension parameters and scopes defined in the OpenID Connect specifications that a Relying Party (OAuth Client) sends to an OpenID Provider (Authorization Server) for the purpose of authenticating an end user.

Authentication Request (SAML 2.0)

A SAML XML document that a Service Provider sends to an Identity Provider to request that the Identity Provider to authenticate the identity of an end user and to return a response for the request.

<AuthnRequest>

See *Authentication Request (SAML 2.0)*

Authorization Request

A request based on the OAuth 2.0 Authorization Framework that an OAuth Client sends to an Authorization Server for the purpose of obtaining an access token (for the purpose of ultimately accessing protected resources on a Resource Server).

attribute source

Specific database or directory location containing data needed by an IdP to fulfill a connection partner's attribute contract or by an SP to look up additional attributes to fulfill an adapter contract.

back-channel

Server-to-server, cross-domain communication path using a protocol, typically SOAP, that does not rely on a browser as an intermediary.

binding

A mapping of SAML request and response messages to specific transport protocols (redirect, POST, or artifact).

certificate

A digital file used for identity verification and other security purposes. The certificate, which is often issued by a Certificate Authority (CA), contains a public key, which can be used to verify the originator's identity.

Certificate Revocation List

(CRL) A list of revoked signing certificates, maintained by the issuing authority at a public URL.

channel

A dedicated Outbound Provisioning configuration specific to a particular service partner, data source, and target service.

connection partner

Entities, such as companies, that are part of an identity federation. These entities are referred to as connection partners in the PingFederate configuration process.

credential

Information used to identify a subject for access purposes (e.g., username and password). A credential can also be a certificate.

database management system

A system for storing and maintaining user account information and attributes. The tables and columns in the RDBMS are used by PingFederate to create user look-up and attribute retrieval queries. (See *Java Database Connectivity*.)

data store

A database or directory location containing user account records and associated user attributes.

Data Encryption Standard (DES)

A symmetric-key method of encryption.

defederation

Optional user-initiated delinking of an identity federation that uses a persistent name identifier or pseudonym for account linking.

digital signature

A process for verifying the identity of the originator of an electronic document and whether the document has been intercepted or altered. The process involves message signing, signature validation, and signing policy coordination between partners.

endpoint

A terminal or gateway that generates or terminates a stream of information. For example, a PingFederate SP server contains an endpoint for the Assertion Consumer URL.

entity ID

The XML element in a SAML assertion that uniquely identifies an Identity Provider.

extensible markup language

A structured, hierarchical text format—based on SGML (Standard Generalized Markup Language)—for the flexible and organized exchange of data.

grant type

The intermediate credentials that represent a resource owner authorization. Grant types are exchanged by the client with the OAuth Authorization Server in order to obtain an access token.

HTTP cookie

Information sent from a server to a web browser to identify a registered website user. Once the cookie is placed in the browser, it is sent back to the server to identify the user every time the user accesses the site. PingFederate's integration adapters interface with the cookie.

HTTP header

The section of an HTTP request or response containing information about the client or the server. PingFederate can use HTTP headers to look up session information passed by the IdP's web application.

HTTP request parameter

A named parameter sent as part of a URL request from a browser to a web server.

ID Token

A security token that contains claims (attributes) about the authentication of an end user (represented by the sub claim in the ID token) by an Authorization Server when using an OAuth Client, and potentially other requested claims. The ID Token is represented as a JSON Web Token (JWT).

identity federation

A trust agreement between or among organizations, implemented using accepted standards, to provide user-authentication tokens and other user or system attributes securely across domains, primarily to enable cross-domain SSO.

Identity Provider

The identity source or SAML authority that authenticates a subject and provides an SP with a security assertion vouching for that authentication.

IdP-initiated SSO or SLO

An identity federation transaction in which the initial action requiring a security context from an IdP occurs at a IdP's site. For example, the user is logged on to the IdP and requests protected resources on an SP. The IdP sends authentication information to the SP.

inbound

A direction of message flow coming into a server relative to the server's identity federation role (IdP or SP). For an IdP, inbound messages include SAML authentication requests. For an SP, inbound messages include SAML assertions.

Java database connectivity (JDBC)

A Java API that allows Java programs to interact with databases.

Kerberos ticket

The security token for the Kerberos protocol.

Key Distribution Center

The control center for authentication and authorization for Kerberos.

keysize

The length (in bits) of each key in a keypair.

keypair

The private key and public key represented by a certificate. PingFederate uses the private key of its keypair(s) to generate signatures for assertions, requests, and responses, as applicable.

Lightweight Directory Access Protocol

A set of protocols used for accessing information directories. PingFederate uses the LDAP v3 protocol for user look-up and attribute processing.

metadata

The SAML 2.0 standards define a metadata exchange schema for conveying XML-formatted information between two SAML entities. Metadata includes endpoint URLs, binding types, attributes, and security-policy information.

Network Access Server

(NAS) A RADIUS client server that provides a single point of access to a protected resource.

OAuth

A standard framework that enables an application (OAuth Client) to obtain access tokens from an OAuth Authorization Server for the purpose of retrieving protected resources on a Resource Server.

OAuth Authorization Server

A server that issues access tokens to clients (sometimes on behalf of a resource owner) for use in authenticating a subsequent Representational State Transfer (REST) API call.

OAuth Client

An application that desires access to a resource protected by a Resource Server and interacts with an Authorization Server to obtain access tokens to do so.

OIDC

See *OpenID Connect*.

Online Certificate Status Protocol

(OCSP) A standard developed by the Internet Engineering Task Force that enables applications to obtain the current status of signing certificates, indicating whether a certificate has been revoked, via HTTP.

opaque

Not readable. If a user's subject identifier is opaque, an SSO partner cannot directly identify the user with reference to the source. An persistent identifier, or *pseudonym*, can be used for Account Linking.

OpenID Connect

OpenID Connect 1.0 is an identity layer on top of the OAuth 2.0 protocol. This protocol enables a Relying Party (OAuth Client) to verify the identity and to obtain claims (attributes) about an end user from an Authorization Server (OpenID Provider).

OpenID Provider

An OAuth 2.0 Authorization Server that is capable of authenticating end users and providing claims on their behalf to Relying Parties.

outbound

A direction of message flow leaving a server. For an IdP, outbound messages include SAML assertions. For an SP, outbound messages include SAML authentication requests.

partner

See *connection partner*.

policy

A set of rules for handling security token requests in PingFederate.

portal

A Web-based application, accessed using a web browser, that often aggregates content from multiple providers, serves as a central point of entry, or both.

POST

An HTTP method of transmitting data contained in HTML forms, by which the data appears in the message body.

Primary Domain Controller

A role that is assigned to a particular server participating in a Windows network.

principal

A user, system, or process whose identity can be authenticated. See *subject*.

profiles

Rules that describe how to embed SAML assertions into and extract them out of other protocols in order to enable SSO or SLO. Profiles describe SAML request and response flows that fulfill specific use cases.

protected resource

Information, typically accessed via a web URL, that is protected by an access management system. See *target URL*.

protocol

An agreed-upon format for transmitting data. XML format of SAML request or response messages.

pseudonym

A persistent name identifier assigned to a user and shared among entities, usually with the user's permission, to enable SSO and SLO. Pseudonyms are often used with the SAML account linking protocol to enable SSO while preventing the discovery of the user's identity or activities.

Public Key Infrastructure

(PKI) Enables users of an unsecured public network, such as the Internet, to securely and privately exchange data and money through the use of keypairs and certificates. The PKI provides for a digital certificate that can identify an individual or an organization and directory services that can store and, when necessary, revoke the certificates.

redirect

A SAML binding that conveys a request or response by sending the user's browser to another location. For instance, an authentication request can be sent from an SP through a browser to an IdP.

refresh token

A long-lived token used by the client to obtain a new access token without having to obtain fresh authorization from the resource owner.

Relying Party

An OAuth 2.0 Client that requires end-user's authenticity and claims (attributes) from an OpenID Provider.

Remote Authentication Dial-in User Service

(RADIUS) A networking protocol for user-access management that includes specifications for two-factor authentication.

<RequestSecurityToken>

(RST) WS-Trust or WS-Federation XML element identifying a request for validation of a security token, or for validation and then issuance of a replacement security token.

<RequestSecurityTokenResponse>

WS-Trust or WS-Federation XML element identifying a response to an RST and containing either the status of the submitted security token or both the status and (if requested and the received token is valid) a newly issued token for further SSO or Web-Services processing.

Resource Server

A server capable of accepting and responding to resource requests on which an access token is presented.

SAML

See *Security Assertion Markup Language*.

SAML authority

A security domain that issues SAML assertions.

scope

Permissions (for example, creating an event on a calendar) associated with an access token.

Secure Sockets Layer

An encryption protocol that sends data between a client and server over a secure HTTP connection.

Security Assertion Markup Language

(SAML) A standard, XML-based, message-exchange framework enabling the secure transmittal of authentication tokens and other user attributes across domains.

System for Cross-domain Identity Management

(SCIM) A REST-based protocol for provisioning and managing user identities across the Internet (see www.simplecloud.info).

security domain

An application or group of applications that trust a common security token used for authentication, authorization, or session management. The token is issued to a user after the user has authenticated to the security domain.

security token

A collection of information used to establish acceptable identity for security purposes. Tokens can be in binary or XML format. A SAML assertion is one kind of security token.

Security Token Service

An entity responsible for responding to WS-Trust requests for validation and issuance of security tokens used for SSO authentication to Web Services.

service-oriented architecture

A loosely coupled application architecture in which all functions or services are accessible via standard protocols. Interfaces are platform and programming-language independent.

Service Provider

A system entity that provides access to a protected resource based on authentication information supplied by an IdP.

SP-initiated SSO or SLO

An identity-federation transaction in which the initial action requiring a security context from an IdP occurs at a SP's site.

session persistence

A mechanism for identifying a user or browser for subsequent requests to a server, needed because the HTTP protocol is stateless. This information is used to lookup state information for the user—for example, items in a shopping cart. PingFederate does not implement session persistence; it facilitates the communication of session information between systems that do.

Simple Object Access Protocol

(SOAP) Defines the use of XML and HTTP to access services, objects, and servers in a platform-independent manner.

single logout

The process of logging a user out of multiple “session participants” or sites where the user has started an SSO session.

Single Logout Return Service

The SAML implementation endpoint URL that returns logout requests.

Single Logout Service

The SAML implementation endpoint URL that receives logout requests for processing.

single sign-on

(SSO) The process of authenticating an identity (signing on) at one website (usually with a user ID and password) and then accessing resources secured by other domains without re-authenticating.

Single Sign-on Service

The SAML implementation endpoint URL that receives authentication requests for processing.

Source ID

A 20-byte sequence used to determine an IdP's identity.

subject

A person, computer system, or application. In the SAML context, assertions make statements about subjects. See *principal*.

target URL

The SP's protected resource; the end destination of an SSO event. See *protected resource*.

transient name identifier

A temporary ID used to preserve user anonymity while facilitating account linking.

token authorization

A mechanism for evaluating attribute criteria available during a transaction to determine whether a user is authorized to access resources. A token in this instance can mean any type of security token—for example, SSO, session cookie, or OAuth token.

token exchange

The process by which a security token is exchanged for another security token.

token translators

An aggregate term for both token processors (used by the IdP PingFederate Security Token Service (STS) to handle different types of incoming security tokens) and token generators (used by the SP PingFederate STS to issue various types of tokens).

Uniform Resource Identifier

Identifies a web resource with a string of characters conforming to a specified format.

Uniform Resource Locator

Identifies a resource according to its Internet location.

virtual server ID

An optional unique identifier by which an identity federation deployment can be known to a specific connection partner.

Web Services Security

A standard mechanism for securing Web Service interactions, often by binding a security token to the Web Service request.

Web Services

Nonbrowser-based, loosely coupled applications that provide modular, programming-language-independent access to specific functions and data across the Internet, via XML and standard protocols.

Web Service Client

An entity that requests a Web Service interaction. In the context of an STS, the Web Service Client would request that a security token be issued for the interaction.

Web Service Enhancement

Supplemental software for the .NET framework provided by Microsoft.

Web Service Provider

In the context of an STS, an entity that requests validation of the security token sent with a client's request for service.

WS-SX

The OASIS committee working on WS-Trust.

WS-Trust

A standard protocol by which an application can request that an STS issue, validate, or exchange security tokens.

List of acronyms

ACS	Assertion Consumer Service
API	Application Programmer Interface
ARS	Artifact Resolution Service
CA	Certificate Authority
CRL	Certificate Revocation List
CSR	Certificate Signing Request
DBMS	Database Management System
DMZ	Demilitarized Zone
DN	Distinguished Name (certificate identifier)
DNS	Domain Name System
EIM	Enterprise Identity Management
GUI	Graphical User Interface
HTTP	HyperText Transfer Protocol
HTTPS	Secure HyperText Transfer Protocol
IdM	Identity Management
IdP	Identity Provider
IP	Internet Protocol
J2SDK	Java 2 Software Development Kit
JDBC	Java Database Connectivity (JDBC)
LDAP	Lightweight Directory Access Protocol
NAS	Network Access Server
O	Organization
OASIS	Organization for the Advancement of Structured Information Standards
OCSP	Online Certificate Status Protocol
OU	Organizational Unit
PKI	Public Key Infrastructure
RADIUS	Remote Authentication Dial-in User Service
RDBMS	Relational Database Management System
RST	
RSTR	

SAML	Security Assertion Markup Language
SaaS	Software as a Service
SCIM	System for Cross-domain Identity Management
SDK	Software Development Kit
SP	Service Provider
SLO	Single Logout
SOA	service-oriented architecture
SOAP	Simple Object Access Protocol
SQL	Structured Query Language
SSL	Secure Sockets Layer
SSL/TLS	Secure Sockets Layer/Transport Level Security
SSO	Single Sign-On
SSTC	Security Services Technical Committee (of OASIS)
STS	Security Token Service
TCP	Transmission Control Protocol
URI	Uniform Resource Identifier
URL	Uniform Resource Locator
WCF	Windows Communication Foundation
WIF	Windows Identity Foundation
WSC	Web Service Client
WSP	Web Service Provider
WSS	Web Services Security
XASP	X.509 Attribute Sharing Profile
XML	Extensible Markup Language


Server clustering guide


PingFederate® provides clustering features that allow a group of PingFederate servers to appear to browsers and partner federation servers as a single system. In this configuration, all client traffic normally goes through a load balancer, which routes requests to the PingFederate servers in the cluster. User-session states and configuration data are shared among the servers, enabling them to process requests as a single entity.

Overview of clustering

PingFederate® provides clustering features that allow a group of PingFederate servers to appear to browsers and partner federation servers as a single system. In this configuration, all client traffic normally goes through a load balancer, which routes requests to the PingFederate servers in the cluster. User-session states and configuration data are shared among the servers, enabling them to process requests as a single entity.

When deployed appropriately, server clustering can facilitate high availability of critical services. Clustering can also increase performance and overall system throughput. It is important to understand, however, that availability and performance are often at opposite ends of the deployment spectrum. Thus, you may need to make some configuration tradeoffs that balance availability with performance to accommodate specific deployment goals. Some of these choices are identified throughout this topic.

 **Important:** All servers in a cluster must use the same version of PingFederate.

 **Note:** PingFederate provides separate failover capabilities specifically for Outbound Provisioning, which by itself does not require either load balancing or state management (see [Deploy provisioning failover](#)).

General architecture

The cluster architecture has two layers. The cluster-protocol layer allows the PingFederate servers to communicate with each other and maintain the cluster as individual servers leave and join. PingFederate's runtime state-management and configuration synchronization services are built on top of the protocol layer. These two layers have different architecture requirements and configuration.

PingFederate abstracts its runtime session-state management behind Java service interfaces. This enables PingFederate to use interface implementations without regard to underlying storage and sharing mechanisms. The abstraction also provides a well-defined point of extensibility in PingFederate (see [Runtime state-management services](#) on page 544).

Group-RPC oriented approach

The prepackaged state-management implementations are designed to accommodate a variety of deployments. The implementations leverage a remote-procedure-call (RPC) framework for reliable group communication, allowing PingFederate servers within a cluster to share state information.

Load balancing

Clustered deployments of PingFederate for single sign-on (SSO) and logout transactions 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 via remote procedure calls as it processes SSO and logout requests.

PingFederate does not automatically balance the traffic among the servers in the cluster. Single sign-on and logout requests must be managed externally to avoid overloading individual servers in the cluster. Because each server can only handle a certain amount of traffic, refer to [PingFederate performance tuning](#) on page 637 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. Specific balancing strategies, their strengths and weaknesses, as well as the impacts on PingFederate are discussed later (see [Runtime state-management architectures](#) on page 539).

Load balancers may 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 (via the administrative console). A cluster may contain one or more runtime nodes but only one console node.



Note: The PingFederate administrative console node must be configured to run outside of the load-balanced group to successfully process SSO requests.

Deploy cluster servers

Follow these steps to configure and deploy clustered PingFederate® servers.



Note: Additional steps are required to set up failover for provisioning. Alternatively, if you are grouping servers *exclusively* to provide for provisioning failover, skip this section and refer to information under [Deploy provisioning failover](#).

1. For each node in a cluster, install PingFederate.
2. For each node, edit its clustering properties in the `run.properties` file located in the `<pf_install>/pingfederate/bin` directory (see [Configure cluster protocol properties](#) on page 535).
3. Optional: For each engine node, edit configuration files that control cluster protocol and runtime session-state management (see [Runtime state-management architectures](#) on page 539 and [Runtime state-management services](#) on page 544).
4. Optional: If Outbound Provisioning is configured at your site and you want to provide failover capabilities, identify and configure the provisioning failover nodes (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, as prompted.

(If you need to request a new license or replace the current one, see [Manage PingFederate license](#) on page 91.)

8. On the **Server Configuration > Cluster Management** screen, click **Replicate Configuration** to push the license information from the console node to all engine nodes.




Note: Starting with PingFederate 8.2, you must use the **Replicate Configuration** workflow 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.

Once 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.



Cluster protocol architecture




PingFederate's cluster protocol services manage cluster messaging, membership, discovery, detection of failed nodes, and merging of split clusters. Nodes in the cluster must be able to communicate with one another over both the cluster bind port and the cluster failure detection port (see [Configure cluster protocol properties](#) on page 535).

-  **Important:** This communication requirement remains true for all cases, including runtime state-management architectures that use designated state servers or subclusters.


Configure cluster protocol properties

The clustering configuration options are maintained in the `run.properties` file located in the `<pf_install>/pingfederate/bin` directory. Configure the clustering properties on each node in the cluster. For information about each property, refer to the following table:


Property	Description
<code>pf.operational.mode</code>	Controls the operational mode of the PingFederate® server. Refer to the properties file for a list of valid values and descriptions. Note that the value <code>STANDALONE</code> should not be used in a cluster unless user-state management is not needed for any reason and configuration-archive deployment is used as the configuration synchronization method.
<code>pf.cluster.node.index</code>	<p>Defines a unique index number for the server in a cluster; the index number is used to identify peers and optimize inter-node communication. (Range: 0 to 65535.)</p> <p>If no value is set for the node index, the system assigns a default index derived from the last two octets of the IP address. We recommend, however, that you assign static indexes.</p> <p> Tip: The PingFederate Cluster Node Selector, available in version 7.0 and higher, provides the capability of overriding adapter processing based on the runtime node servicing a request.</p>
<code>pf.cluster.auth.pwd</code>	<p>Sets the password that each node in the cluster must use to authenticate when joining the group. This prevents unauthorized nodes from joining a cluster. (Value: any string, or blank.)</p> <p>Consider using a randomly-generated key with 22 or more alphanumeric characters as the property value. While optional, we recommend that you obfuscate the property value. For information about the <code>obfuscate</code> command-line utility, see its built-in help.</p> <p>All nodes in a cluster must share the same property value, blank or otherwise.</p>
<code>pf.cluster.encrypt</code>	<p>Indicates whether or not to encrypt network traffic sent between nodes in a cluster. (Values: <code>true</code> <code>false</code>.) The default value is <code>false</code>.</p> <p>When set to <code>true</code>, communication within the cluster is encrypted with a symmetric key derived from the value of the <code>pf.cluster.auth.pwd</code> property.</p> <p> Important: When the <code>pf.cluster.encrypt</code> property is set to <code>true</code>, you must provide a value for the <code>pf.cluster.auth.pwd</code> property; otherwise PingFederate aborts during its startup process.</p> <p>All nodes in a cluster must have the same value set for this property.</p>
<code>pf.cluster.encryption.keysize</code>	<p>The length of the key that PingFederate takes into consideration when deriving the symmetric key from the value of the <code>pf.cluster.auth.pwd</code> property for the purpose of encrypting network traffic sent between nodes in a cluster. Applicable only when the <code>pf.cluster.encrypt</code> is set to <code>true</code>. The default value is 128.</p> <p>All nodes in a cluster must have the same value set for this property.</p>
<code>pf.cluster.bind.address</code>	Defaults to <code>NON_LOOPBACK</code> , which leaves the system to choose an available non-loopback IP address. Alternatively, enter an IP address of the network interface to which the group communication should bind. For machines with more than one network interface, provide a specific IP address. You can use this property to increase

Property	Description
	<p>performance (particularly with UDP) and improve security by segmenting group-communication traffic onto a private network or VLAN.</p> <p> Tip: Besides <code>NON_LOOPBACK</code> or an IP address, you can also use other values supported by JGroups. For more information, search for the parameter <code>bind_addr</code> in JGroups documentation (jgroups.org/manual/index.html#Transport).</p> <p> Important: This field does not support DNS name. Use the default value (<code>NON_LOOPBACK</code>) or replace it with an IP address.</p>
<code>pf.cluster.bind.port</code>	<p>Specifies the port associated with the <code>pf.cluster.bind.address</code> property or with the default network interface used. The default value is <code>7600</code>.</p> <p>This is the port used by other cluster members during their discovery process (usually via the <code>pf.cluster.tcp.discovery.initial.hosts</code> property).</p>
<code>pf.cluster.failure.detection.bind.port</code>	<p>Indicates the bind port of a server socket that is opened on the given node and used by other nodes as part of one of the cluster's failure-detection mechanisms. If zero or unspecified, a random available port is used. The default value is <code>7700</code>.</p>
<code>pf.cluster.transport.protocol</code>	<p>Indicates the transport protocol used for group communication. (Values: <code>udp</code> <code>tcp</code>.) The default value is <code>tcp</code>.</p> <p>Use UDP multicast when IP multicasting is enabled in the network environment and the majority of group traffic is point-to-full-group. You must also provide values for the <code>pf.cluster.mcast*</code> properties described below.</p> <p>Use TCP for geographically dispersed servers or when multicast is not available or disabled for some other reason (routers discard 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 provide a value for the <code>pf.cluster.tcp.discovery.inital.hosts</code> property described below.</p> <p> Note: This property is a reference to a protocol-stack XML configuration file located in the <code><pf_install>/pingfederate/server/default/conf/</code> directory. Two stacks are provided: one for UDP multicast and one for TCP. Administrators can customize either stack or add to it as needed by modifying the associated configuration file.</p>
<code>pf.cluster.mcast.group.address</code>	<p>Defines the IP address shared among nodes in the same cluster for UDP multicast communication; required when UDP is set as the transport protocol. (Range: <code>224.0.0.0</code> to <code>239.255.255.255</code>; note that some addresses in this range are reserved for other purposes.) This property is not used for TCP.</p> <p>All nodes in a cluster must use the same address for this property and the port property below.</p>
<code>pf.cluster.mcast.group.port</code>	<p>Defines the port associated with the property above, <code>pf.cluster.mcast.group.address</code>.</p>
<code>pf.cluster.tcp.discovery.initial.hosts</code>	<p>Designates a static list of PingFederate servers to be contacted for group membership information when discovering, joining, and rejoining the group; 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:</p> <pre>host1[7600],10.0.1.4[7600],host7[1033],10.0.9.45[2231]</pre> <p>When using static discovery, this list should contain all nodes. It is recommended to include the current node in the list, so each node should have an identical value for this property.</p>

Property	Description
	Alternatively, leave this property blank and enable dynamic discovery in the <code><pf_install>/pingfederate/server/default/conf/tcp.xml</code> file.
<code>pf.cluster.diagnostics.enabled</code>	Indicates if JGroups diagnostics is turned off (false) or on (true). The default value is false.
<code>pf.cluster.diagnostics.addr</code> and <code>pf.cluster.diagnostics.port</code>	The multicast address and port this node listens on for diagnostic messages. Do not change the property values.

 **Note:** This table does not include information about properties used for provisioning failover (see [Deploy provisioning failover](#)).

1. Configure the clustering properties for one of the nodes in the cluster.
2. Repeat this process on the rest of the nodes in the cluster.


 **Important:** You must manually configure the clustering properties in the `run.properties` file on each node. The `run.properties` file is *not* copied from the console node to the engine nodes automatically; it is also *not* part of the **Replicate Configuration** process. PingFederate servers must be restarted if running.

Enable dynamic discovery for clustering


Dynamic discovery is well suited for environments where traffic volume may spike and require additional resources during the peak period to handle the increased traffic. The elastic scaling capability helps you to bring additional PingFederate® engine nodes online with no additional configuration changes after the initial setup.

When this discovery mechanism is enabled, a new group member pulls cluster membership information from a centralized repository. If a cluster exists, it joins the group; otherwise, it forms a new group. It is crucial that the information is safely stored and readily accessible by all nodes. PingFederate supports IAM roles for Amazon Elastic Compute Cloud (Amazon EC2), Amazon Simple Storage Service (Amazon S3), and OpenStack Swift, giving you the flexibility between public and private cloud storage for the dynamic discovery mechanism.

Dynamic discovery requires only a one-time setup. Its configuration is maintained in the `tcp.xml` file located in the `<pf_install>/pingfederate/server/default/conf` directory. Once configured, you are no longer required to coordinate the effort of updating IP addresses on all nodes.

 **Important:** You must manually configure or synchronize the dynamic discovery properties in the `tcp.xml` file on each new node. The `tcp.xml` file is *not* synchronized automatically across the nodes in a cluster; it is also *not* part of the **Replicate Configuration** process. PingFederate servers must be restarted if running.

1. Configure all required `pf.cluster.*` properties in the `run.properties` file in the `<pf_install>/pingfederate/bin` directory.

 **Note:** Ensure the `pf.cluster.transport.protocol` property is set to `tcp` (the default value).

2. Edit the `tcp.xml` file in the `<pf_install>/pingfederate/server/default/conf` directory.

Follow the inline comments to enable dynamic discovery.


Field	Description
IAM roles for Amazon EC2 (<code>com.pingidentity.aws.AWS_PING</code>)	
<code>port_number</code>	The port, on which PingFederate listens for cluster communication. The default value is <code>\${pf.cluster.bind.port}</code> , which pulls the value defined in the <code><pf_install>/pingfederate/bin/run.properties</code> file.
<code>tags</code>	A comma separated list of EC2 tags.


Field	Description
	<p>When specified, only Amazon Machine Images (AMIs) that have been assigned with the specified tag (or tags) are eligible to join the cluster. If multiple tags are specified, only AMIs that have been assigned with all tags will be considered.</p> <p>For information about tags, see the documentation from Amazon EC2 (docs.aws.amazon.com/AWSEC2/latest/UserGuide/Using_Tags.html).</p>
filters	<p>A semi-colon separated list of key value pairs of system metadata.</p> <p>When specified, only AMIs that match the specified filter (or filters) are eligible to join the cluster. If multiple filters are specified, only AMIs that match all filters will be considered.</p> <p>Note that you can enter a comma separated list of values for each filter. In this case, a filter is considered a match so long as one of the values is satisfied. For example, if you enter:</p> <pre>filters="instance-type=t2.small,t2.medium;architecture=x86_64"</pre> <p>then only AMIs that have an instance type of either t2.small or t2.medium and also an architecture of x86_64 will be considered.</p> <p>For information about filters, see the documentation from Amazon EC2 (docs.aws.amazon.com/cli/latest/reference/ec2/describe-instances.html).</p> <p>You can use a combination of tags and filters, in which case only AMIs that satisfy both criteria are eligible to join the cluster.</p>
access_key and secret_key	<p>The access key and its secret key for the purpose of querying EC2 for AMIs.</p> <p>If access_key and secret_key are defined, tags and filters are ignored.</p> <p>Keys can be encrypted using the obfuscate utility (obfuscate.bat for Windows or obfuscate.sh for Unix or Linux), located in the <pf_install>/pingfederate/bin directory.</p>
log_aws_error_messages	<p>When set to true (the default), error messages received from Amazon EC2 are logged to the server log.</p>
Amazon S3 (S3_PING)	
location	<p>The name of the bucket in your Amazon S3 environment.</p> <p>For information about buckets, see the documentation from Amazon Web Services (docs.aws.amazon.com/AmazonS3/latest/gsg/CreatingABucket.html).</p>
access_key and secret_access_key	<p>The security credentials to access your Amazon S3 environment.</p> <p>For information about both keys, see the documentation from Amazon Web Services (docs.aws.amazon.com/general/latest/gr/aws-sec-cred-types.html#access-keys-and-secret-access-keys).</p> <p>Keys can be encrypted using the obfuscate utility (obfuscate.bat for Windows or obfuscate.sh for Unix or Linux), located in the <pf_install>/pingfederate/bin directory.</p>
remove_all_files_on_view_change	<p>When set to true (the default), JGroups cleans up files when it detects a view change. (For more information, see http://jgroups.org/manual/index.html#FILE_PING.)</p>

OpenStack Swift (SWIFT_PING)

Field	Description
auth_type	The authentication type.
auth_url	The authentication URL.
username and password	The security credentials. Password can be encrypted using the obfuscate utility (<code>obfuscate.bat</code> for Windows or <code>obfuscate.sh</code> for Unix or Linux), located in the <code><pf_install>/pingfederate/bin</code> 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 fields, see http://jgroups.org/manual/index.html#_swift_ping .)
remove_all_files_on_view_change	When set to <code>true</code> (the default), JGroups cleans up files when it detects a view change. (For more information, see http://jgroups.org/manual/index.html#FILE_PING .)

- Restart PingFederate if it is running.
- Repeat these steps on the rest of the nodes (and any new nodes, as needed).

 **Tip:** After the initial setup, your nodes are ready to be deployed, undeployed, and redeployed as traffic volume changes.

 **Note:** The configuration and process of the discovery mechanism are independent from those of the runtime state-management services.

The goal of the discovery mechanism is to find a node to retrieve cluster information and to subsequently joins and rejoins such cluster while the runtime state-management services are about sharing session-state information with the applicable nodes based on the associated configurations. Some services support the concept of subcluster; however, any node in a subcluster must still be able to communicate with other nodes in the cluster. For additional information, see [Runtime state-management architectures](#) on page 539.

Runtime state-management architectures

In a cluster consisting of a large number of nodes, it may be desirable to share session-state information with only a subset of nodes (for performance reasons). To facilitate this, most of the group RPC-based service implementations (see [Runtime state-management services](#) on page 544) make use of a *preferred-nodes* concept, which allows each node to have 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. Any changes must be replicated manually for each cluster node. The table below indicates the configuration file that applies to each implementation. Refer to the indicated sections in this topic for detailed information about each implementation.

Configuration File	RPC-Based Service Implementation
cluster-account-locking.conf	Account Locking Service on page 549
cluster-artifact.conf	Artifact-Message Persistence and Retrieval Service on page 547
cluster-assertion-replay-prevention.conf	Assertion Replay Prevention Service on page 546
cluster-idp-session-registry.conf	IdP Session Registry Service on page 546

Configuration File	RPC-Based Service Implementation
cluster-inter-request-state.conf	Inter-Request State-Management (IRSM) Service on page 544
cluster-session-revocation.conf	Back-Channel Session Revocation Service on page 548
cluster-sp-session-registry.conf	SP Session Registry Service on page 546

The following table describes properties contained in the configuration files (the Artifact-Message Persistence and Retrieval Service uses only the `rpc.timeout` setting):

Property	Description
preferred.node.indices	<p>A comma-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.</p> <p>The Artifact-Message Persistence and Retrieval Service and the Back-Channel Session Revocation Service do not support this parameter.</p>
preferred.node.group.id (found only in cluster-idp-session-registry.conf)	<p>An alphanumeric group ID for each subcluster. If specified, the group ID must be unique for each subcluster. 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.</p> <p>In addition, 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 in the IdP Configuration (or SP Configuration) > Sessions screen. 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.</p> <p>If subclusters are configured without specifying any 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.</p>
rpc.timeout	How long, in milliseconds, this node waits before timing out unresponsive RPC invocations.
synchronous.retrieve.majority	<p>Indicates how many responses to wait for when making synchronous remote procedure calls (values: <code>true</code> <code>false</code>). When set to <code>true</code> (the default), this node waits for the majority of recipients to respond. When set to <code>false</code>, it waits for all recipients to respond.</p> <p>Note this property is not applicable to the Account Locking Service and not found in the <code>cluster-account-locking.conf</code> file.</p>
bulk.revoked.sris.timeout (found only in cluster-session-revocation.conf)	A node downloads a full revocation list from another node during startup or when it rejoins a cluster after being disconnected from it (possibly 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 milliseconds (10 seconds).
read.local.only (found only in cluster-session-revocation.conf)	Determines whether PingFederate should process queries for revocation status by searching the local revocation list or collecting the information from other engine nodes in the cluster.

Property	Description
	<p>When set to <code>true</code> (the default), queries for revocation status are processed locally. When <code>false</code>, the processing node pulls revocation status from other engine nodes in the cluster (subject to the <code>rpc.timeout</code> value).</p> <p> 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 548).</p>


Preferred node indices

Configuring the `preferred.node.indices` property could reduce the memory footprint and network communications; however, care must be given as results vary depending on the volume of transactions and the distribution of them across engine nodes. Performance tuning could help in this regard. For more information, see [PingFederate performance tuning](#) on page 637.

Additionally, note that within a single cluster deployment, individual services can use different preferred nodes. In other words, you can set different values for the `preferred.node.indices` property for each service.

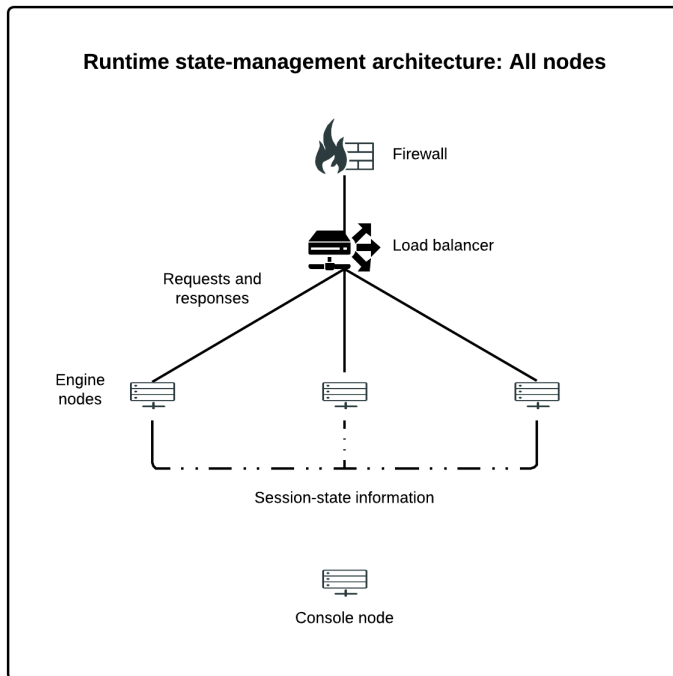
The use of preferred nodes can translate into any number of deployment configurations. Three primary strategies are discussed in the following sections and may serve as touch points for you to consider in conjunction with your network requirements:

- [Sharing all nodes](#) on page 541
- [Designating state servers](#) on page 542
- [Defining subclusters](#) on page 543

 **Note:** For PingFederate versions 7.0 and higher, it is possible to configure overrides for authentication-adaptor processing based on the runtime node servicing a request (see [Configure the Cluster Node Authentication Selector](#)).

Sharing all nodes

Leaving the `preferred.node.indices` property blank in all the cluster-configuration files provides 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.



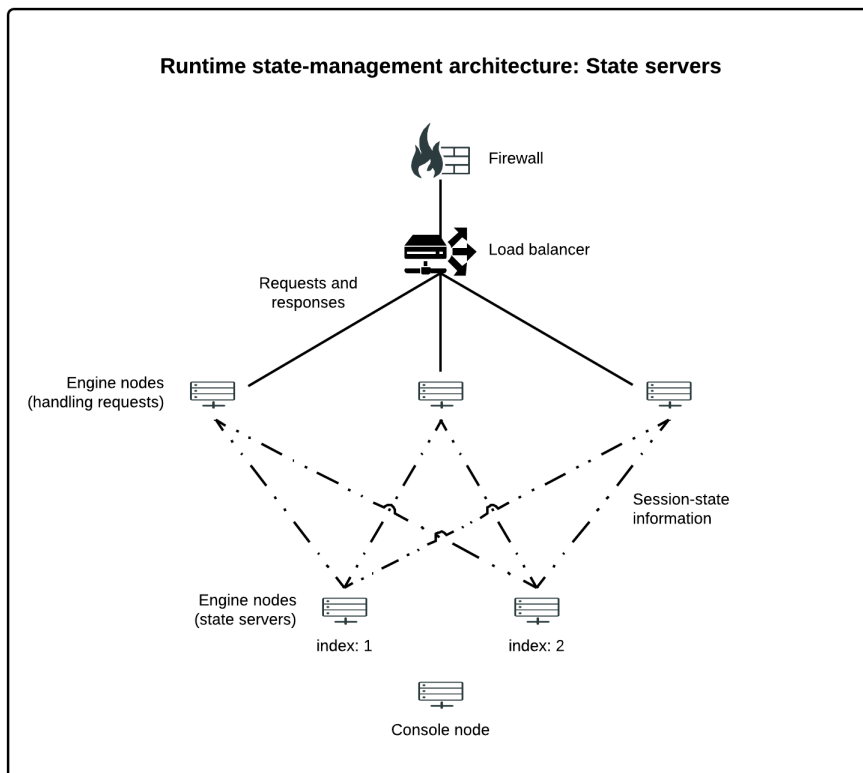
The diagram above illustrates the node-sharing approach. HTTP traffic is directed to all nodes.

Designating state servers

You can select a few nodes as *state servers*. This deployment can be configured by setting the preferred-node indices of other servers in a group to those of the state servers. The load balancer should be configured to isolate the state-server nodes from end-user traffic. This approach scales better than the all-nodes approach because additional nodes do not require as much communication to every other node.



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.



This diagram illustrates the state-server approach. If the two state-server nodes have indices of 1 and 2, then the preferred-nodes property of the runtime 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.



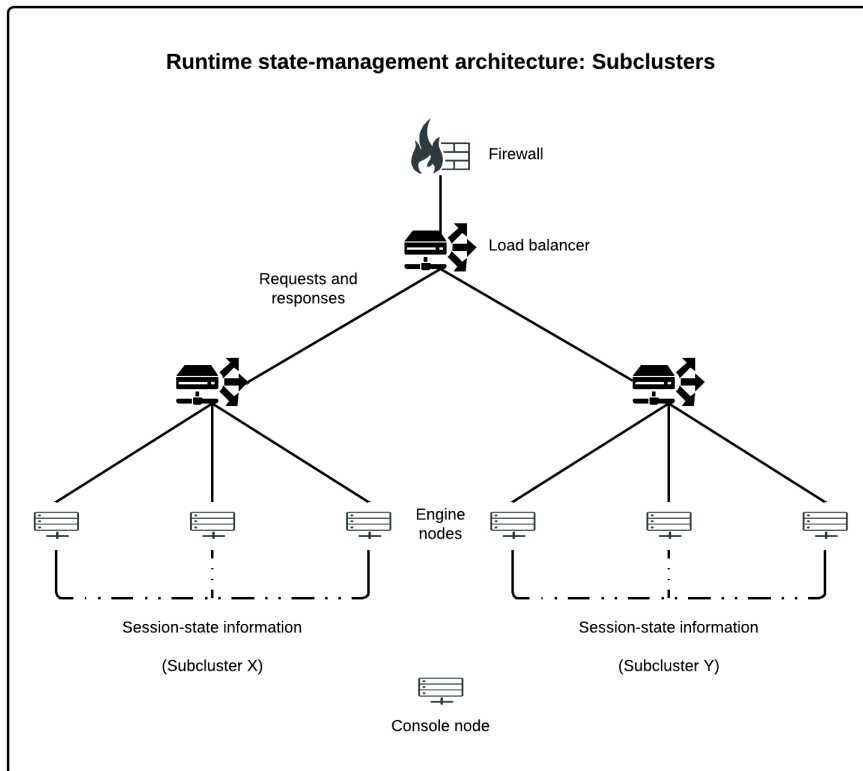
Note: When PingFederate acts as an OAuth authorization server (AS) and the access token manager uses a reference token data model, the resource server (RS) must send a request to PingFederate to de-reference the access token for the corresponding identity and security information. Because the OAuth clients and the RS send their requests separately, reference token information is shared among all nodes despite any state server or subcluster setup.

Defining subclusters

You can use node indices 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 the load balancer to persist, or *stick*, user sessions to the same group so that each subsequent request from the same user is directed to one of the nodes in that group. (A wide variety of persistence mechanisms exist and may vary by load balancer – consult the documentation specific to your product.) 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.



The diagram above illustrates the subcluster approach. The preferred-nodes property of each server in the cluster lists the indices of all nodes in its subgroup (including itself). HTTP traffic is 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 manager 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, reference token information is shared among all nodes despite any state server or subcluster setup.

Runtime state-management services

The runtime state-management services are a collection of interfaces defining the contract that PingFederate® uses with each service to manage session states. The implementation of each service interface is specified in the `hivemodule.xml` file in the `<pf_install>/pingfederate/server/default/conf/META-INF/` directory. This file does not need to be modified unless you wish to customize the way services are handled in a cluster. Any changes must be replicated manually for 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, if the server is in standalone mode, an in-memory-only implementation is used. If the server is in a clustered mode, a group RPC-based implementation is used. The proxies are provided for convenience; if you choose, you may specifically designate the implementation you desire for each service, as described in the following sections. Configuration files for the implementations that make use of preferred nodes are in the directory `<pf_install>/pingfederate/server/default/conf` (see [Runtime state-management architectures](#) on page 539).

Inter-Request State-Management (IRSM) Service

The PingFederate® server tracks user-session state information between HTTP requests – for example, 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 redirection. Another example is keeping track of state between issuing a request to a partner and processing the response. Generally, this state is short-lived.

The `InterRequestStateMgmtProxy` interface chooses from among three methods available to track this state: group RPC-based (the clustering default), cookie based, and local memory-based.

 **Caution:** Depending on deployments, if cookies are used for state management in a cluster instead of the default RPC-based method, keep in mind that the cookies may become large enough to exceed certain browser limitations.

Group RPC-based session tracking


This implementation uses the preferred-nodes concept. `<pf_install>/pingfederate/server/default/conf/cluster-inter-request-state.conf` is the configuration file. All three preferred-node approaches described under [Runtime state-management architectures](#) on page 539 are possible with this implementation. The subgroup approach, however, may be most appropriate.

The service proxy `InterRequestStateMgmtProxy` in the `hivemodule.xml` file is set to use this implementation as the clustering default. The specific class name is:

```
org.sourceid.saml20.service.impl.grouprpc.InterRequestStateMgmtGroupRpcImpl
```

Cookie-based tracking

In a clustered mode, this implementation tracks short-lived session-state data with the help of browser cookies. The implementation works by compressing and encrypting the session-state data and sending it to the user's browser as a cookie. User data is then returned to the server on subsequent requests from the browser. No data for this service needs to be shared among nodes, and even the simplest round-robin load-balancing technique works. There is some overhead in the encryption and compression, but the advantage is that the session state must travel over the network only between the entities that actually need it – the user and the server that handles that user's next request.

 **Caution:** Cookie-based tracking, which can result in larger cookies, should not be used for configurations using account linking. (For more information, see [About account linking](#).)

This implementation does not use group RPC to share session-state data. However, RPC is used by default to handle dynamic distribution of the AES keys used for encryption and decryption (unless a static key is configured). The `key-tracker.xml` file, in the `<pf_install>/pingfederate/server/default/data/config-store` directory, controls the interval at which new keys are issued, the maximum number of keys to retain, and the key length. The oldest member of a cluster group is responsible for issuing and distributing new keys.


The file `inter-req-cookie-config.xml`, also in the `config-store` directory, specifies the name of the cookie, which you may modify as needed. You can also configure a static AES encryption key by specifying its hex value (the key must be set on each node). While a static key is somewhat less secure, it eliminates the need for key management between nodes via back-channel communication.

To use this service implementation, set the class attribute for the `InterRequestStateMgmt` service in the `hivemodule.xml` file to the class name:

```
org.sourceid.saml20.service.impl.cookie.InterReqStateMgmtCookieImpl
```

Local memory-based tracking

In this alternative, the inter-request state of a user is tracked in the local memory of the processing server.

 **Note:** To implement this alternative, the load balancer must support sticky sessions to force all requests for the same user session to be routed to the same server.

To use this service implementation, set the class attribute for the `InterRequestStateMgmt` service in the `hivemodule.xml` file to the class name:

```
org.sourceid.saml20.service.impl.localmemory.InterReqStateMgmtMapImpl
```

IdP Session Registry Service

PingFederate® uses the IdP Session Registry Service to facilitate single logout by tracking assertions issued to SP partners. This service is used only when the PingFederate server is acting in an IdP role and supports single logout 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 `cluster-idp-session-registry.conf` in the `<pf_install>/pingfederate/server/default/conf` directory.

All preferred-node approaches are possible with this implementation (see [Runtime state-management architectures](#) on page 539). However, if your site accepts single logout messages via SOAP, the subgroup approach may not succeed because the load balancer's sticky-session functionality cannot ensure that the SOAP message from a partner will be directed to the appropriate subgroup.

The service proxy uses the class:

```
org.sourceid.saml20.service.impl.grouprpc.IdpSessionRegistryGroupRpcImpl
```

SP Session Registry Service

PingFederate® uses the SP (Service Provider) session registry service to facilitate single logout by tracking assertions issued from IdP partners. This service is used only when the PingFederate server is acting in an SP role and supports single logout 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 `cluster-sp-session-registry.conf` in the `<pf_install>/pingfederate/server/default/conf` directory.

All preferred-node approaches are possible with this implementation (see [Runtime state-management architectures](#) on page 539). However, if your site accepts single logout messages via SOAP, the subgroup approach may not succeed because the load balancer's sticky-session functionality cannot ensure that the SOAP message from a partner will be directed to the appropriate subgroup.

The service proxy uses the class:

```
org.sourceid.saml20.service.impl.grouprpc.SpSessionRegistryGroupRpcImpl
```

LRU memory management schemes

During the normal course of transaction processing, the Inter-Request State-Management Service and the IdP and SP session registries manage memory as part of normal processing. However, it is common for end users to abandon web sessions, resulting in orphaned data. To ensure that such 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. When a data structure reaches the maximum size, the oldest entries are automatically removed.

The maximum size of each data structure is configurable in the file `size-limits.conf` in the `<pf_install>/pingfederate/server/default/conf` directory.

Assertion Replay Prevention Service

The SAML (Security Assertion Markup Language) standards specify that when an SP receives assertions via the POST binding, the SP should keep track of that assertion for the duration of its validity to ensure that it is not replayed (that is, intercepted by a third party and re-posted). PingFederate® delegates that responsibility to the Assertion Replay Prevention Service. This service is used only when PingFederate is acting as an SP and receives assertions via the POST binding (see various binding profiles under [SAML 1.x profiles](#) and [SAML 2.0 profiles](#)).

When PingFederate is in clustered mode, the service proxy uses a group RPC-based, preferred-nodes implementation; the configuration file is `cluster-assertion-replay-prevention.conf` in the `<pf_install>/pingfederate/server/default/conf` directory.

Due to the nature of the threat that this service is intended to prevent, only the [Sharing all nodes](#) on page 541 and [Designating state servers](#) on page 542 deployment strategies are appropriate for this service (see after [Runtime state-management architectures](#) on page 539).

The service proxy uses the class:

```
org.sourceid.saml20.service.impl.grouprpc.AssertionReplayPreventionServiceGroupRpcImpl
```

This service is unique in that it fulfills only a security condition of the federation specifications, rather than supporting normal SSO functionality. Because many other security requirements are in place (for example, SSL, no-cache directives, and digital signatures), there may be situations where the priority placed on cluster performance outweighs the priority placed on this security check. If you are in this situation, you may wish to change the implementation for the service point `AssertionReplayPreventionService` in the `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; however, you may wish to use it, with caution, when performance is an absolute priority.

Artifact-Message Persistence and Retrieval Service

When the artifact binding is used to send messages to partners, the PingFederate® server must keep track of the message referenced by the artifact until the partner makes a SOAP call to retrieve it.

Two options are available for using outbound artifacts in a cluster:

- Group RPC-based retrieval
- Using SAML 2.0 indexing

Group RPC-based retrieval

When PingFederate is in clustered mode, the service proxy selects a group RPC-based implementation, which takes advantage of node indexing but does not use the preferred-nodes concept. Sticky-session load-balancing strategies are not effective for the artifact binding, because the requests that result in the issuance of the artifact and the artifact resolution request come from different locations.

This implementation works by having the node that issues an artifact encode its node index into the artifact using two bytes of the artifact message handle. When a server receives the artifact resolution request, it uses the encoded node index to determine the issuing node and makes an RPC call to get the associated message.

Although this implementation does not use the preferred-nodes concept, it does have a configuration file, `<pf_install>/pingfederate/server/default/conf/cluster-artifact.conf`, in which the RPC time-out can be configured.


The service proxy uses the class:

```
org.sourceid.saml20.service.impl.grouprpc.ArtifactPersistenceSvcGroupRpcEncodedNodeIdxImpl
```

SAML 2.0 indexing (local memory)

Due to the implementation complexity involved in managing state with the artifact binding, the SAML 2.0 specification introduced the concept of indexed endpoints. A SAML 2.0 federation entity may 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.

The downside to this approach is that partners must know about each of your back-end servers. Generally, this means providing partners with a list that includes multiple artifact-resolution service endpoints with the corresponding indices.

 **Note:** 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/
ARS.ssaml2" index="1"/>
<ArtifactResolutionService Binding="..." Location="https://node2/idp/
ARS.ssaml2" index="2"/>
<ArtifactResolutionService Binding="..." Location="https://node3/idp/
ARS.ssaml2" index="3"/>
<ArtifactResolutionService Binding="..." Location="https://node4/idp/
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 zero in standalone mode).

Not only do partners need to know about each back-end server, they need direct access to each ARS endpoint. This may require more complicated configuration of load balancers, proxies, and firewalls. Another drawback to this approach is that it 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 `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 query the revocation status (see [Session Revocation API](#)).

When PingFederate is in clustered mode, the service proxy uses a group 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. It does not use the preferred-nodes concept. This enables the flexibility of allowing the queries to be processed locally or results to be returned after collecting the information from other engine nodes; the former yields faster response time for engine nodes that are deployed in well-connected networks while the latter adds a layer of protection against inconsistent revocation lists among the engine nodes due to possible network outages.

The configuration file is `<pf_install>/pingfederate/server/default/conf/cluster-session-revocation.conf`. This is where the RPC time-out and other settings can be tuned (see [Runtime state-management architectures](#) on page 539).

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, in addition to the Session Revocation Lifetime setting (see [Configure OAuth AS settings](#)), 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 `<pf_install>/pingfederate/server/default/conf/size-limits.conf`. The default value is 50000.

Account Locking Service

Account lockout protection prevents user accounts from becoming locked 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®; for example, the HTML Form Adapter and the Username Token Processor. (For configuration information, see [Account lockout protection](#)). Account lockout protection is also shared across nodes, which helps in scenarios where PingFederate is deployed behind a load balancing infrastructure without sticky sessions.

When PingFederate is in clustered mode, the service proxy uses a group RPC-based, preferred-nodes implementation; the configuration file is `cluster-account-locking.conf` in the `<pf_install>/pingfederate/server/default/conf` directory.

Extensibility and other services

Custom implementations of all state management services described in this *Guide* can be plugged into PingFederate® as needed for special deployments. Software developers can refer to the Javadoc reference:

```
<pf_install>/pingfederate/sdk/doc/index.html
```

Two other services may need consideration when running PingFederate in a cluster, depending on SAML 2.0 federation deployment needs:

- Account linking service
- Pseudonym service

Account linking service

This service stores the association between the external and internal identifiers of an end user when account linking is used as an SP identity-mapping strategy. The default, standalone implementation uses a JDBC interface to an embedded database within PingFederate. No information from the embedded database is shared across the cluster. Therefore, when account linking is used for an IdP connection deployed in a cluster, the default implementation will not work properly. In such cases, the pointer must be adjusted for cluster use by pointing the service to an external database (see [Define an account-linking data store](#)).


Pseudonym service

This service references the method needed by PingFederate to generate or look up a pseudonym for a user. The service is used 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. However, it may be desirable in some situations to implement pseudonym handling differently. Developers can refer to the Javadoc reference describing `PseudonymService` interface for more information.

Deploy provisioning failover



After configuring outbound provisioning, you have the option to set up one or more failover PingFederate® servers specifically for provisioning backup.


Provisioning runtime processing and failover is independent of SSO or 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.

 **Important:** The pre-installed HyperSQL database cannot be used in a failover configuration. Each server in the failover network must be configured to use the same relational database.

1. Identify two or more runtime instances of PingFederate to configure for provisioning failover.

- For each server instance, edit provisioning properties in the `<pf_install>/pingfederate/bin/run.properties` file as follows:

Property	Description
<code>pf.provisioner.mode</code>	<p>The status of outbound provisioning. Allowed values are:</p> <ul style="list-style-type: none"> <code>OFF</code> – Outbound Provisioning is disabled (the default). <code>STANDALONE</code> – Provisioning is enabled, without failover. <code>FAILOVER</code> – Provisioning is enabled, with failover. <p> Important: The <code>STANDALONE</code> setting is not used for failover configuration; primary and secondary servers must be set to <code>FAILOVER</code>.</p>
<code>provisioner.node.id</code>	<p>The unique index number of the provisioning server.</p> <p>Each server must have a unique index number (from 1 to n), which is used to prioritize which server is currently active and which is next in line in case of a failure.</p> <p> Important: The active primary server <i>must</i> have an index number of 1.</p>
<code>provisioner.failover.grace.period</code>	<p>The 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 Server Configuration > Server Settings > Outbound Provisioning screen on the administrative console.</p>

-  **Important:** You must manually 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.


- Start or restart all of the PingFederate servers.
- If you have not already done so, set up an external database to facilitate provisioning and then update the **Internal Provisioning Data Store** setting in the **Server Configuration > Server Settings > Outbound Provisioning** screen.

Once configured, if the provisioning servers belong to the same PingFederate clustered environment, go to the **Server Configuration > Cluster Management** screen and 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 data store 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 via 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 (for example, in `<pf_install>/pingfederate/bin/run.properties`) must be replicated manually across the cluster, as applicable, and PingFederate servers must be restarted if running.

Console configuration push

When running in cluster mode, the administrative console provides a **Server Configuration > Cluster Management** screen.



Whenever applicable changes are made through the administrative console, a message appears at top of the console to serve as a reminder to go to the **Cluster Management** screen and to replicate the current console configuration to all server nodes in the cluster.

Note: Starting with PingFederate® 8.2, administrators must also use the **Replicate Configuration** workflow to initiate the transmission of the license file from the console node to all server nodes.

The replication process takes only a moment, during which time any new requests coming into any of the cluster nodes are held in a queue. When the replication is complete, processing continues where it left off.

Administrators may also click **Refresh Table** to verify that any new nodes have joined the cluster.

Configuration-archive deployment

After you configure or reconfigure the console, you can also update cluster nodes by downloading a configuration archive from the **Server Configuration > Configuration Archive** screen and then deploying it (either manually or via a scripted process) to the `<pf_install>/pingfederate/server/default/data/archive` 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 the previous section. However, this option provides for scheduling and scripting cluster synchronization.

Runtime state-management services

If you have configured one of the following runtime state-management services on the engine nodes, you must migrate the configuration files (located in the `<pf_install>/pingfederate/server/default/conf` directory) to the engine nodes manually.

Configuration File	RPC-Based Service Implementation
cluster-account-locking.conf	Account Locking Service on page 549
cluster-artifact.conf	Artifact-Message Persistence and Retrieval Service on page 547
cluster-assertion-replay-prevention.conf	Assertion Replay Prevention Service on page 546
cluster-idp-session-registry.conf	IdP Session Registry Service on page 546
cluster-inter-request-state.conf	Inter-Request State-Management (IRSM) Service on page 544
cluster-session-revocation.conf	Back-Channel Session Revocation Service on page 548
cluster-sp-session-registry.conf	SP Session Registry Service on page 546

SSO integration overview

Ping Identity®'s PingFederate® must be integrated programmatically with end-user applications and identity management (IdM) systems to complete the "first- and last-mile" implementation of a federated-identity network. The purpose of this document is to provide an overview of the various 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, PingFederate provides commercial integration kits, which include *adapters* that plug into the PingFederate server and *agents* that interface with local IdM systems or applications.

- [Integration introduction](#) on page 552
- [SSO integration concepts](#) on page 552
- [Identity provider integration](#) on page 553
- [Service provider integration](#) on page 555
- [Summary](#) on page 557

Integration introduction

As a stand-alone server, PingFederate® must be integrated programmatically with end-user applications and identity management (IdM) systems to complete the “first- and last-mile” implementation of a federated-identity network. The purpose of this document is to provide an overview of the various 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, PingFederate provides commercial integration kits, which include *adapters* that plug into the PingFederate server and *agents* that interface with local IdM systems or applications.

Integration kits, including various connectors for secure SSO to Software-as-a-Service (SaaS) providers, are available for download at pingidentity.com. User guides for the kits can be found in the Ping Identity [documentation](#) website..

PingFederate also includes a robust software development kit (SDK), which software developers can use to write their own custom interfaces for specific systems. Please refer to the PingFederate [SDK developer's guide on page 559](#) for more information, available in the PingFederate distribution `sdk` directory.

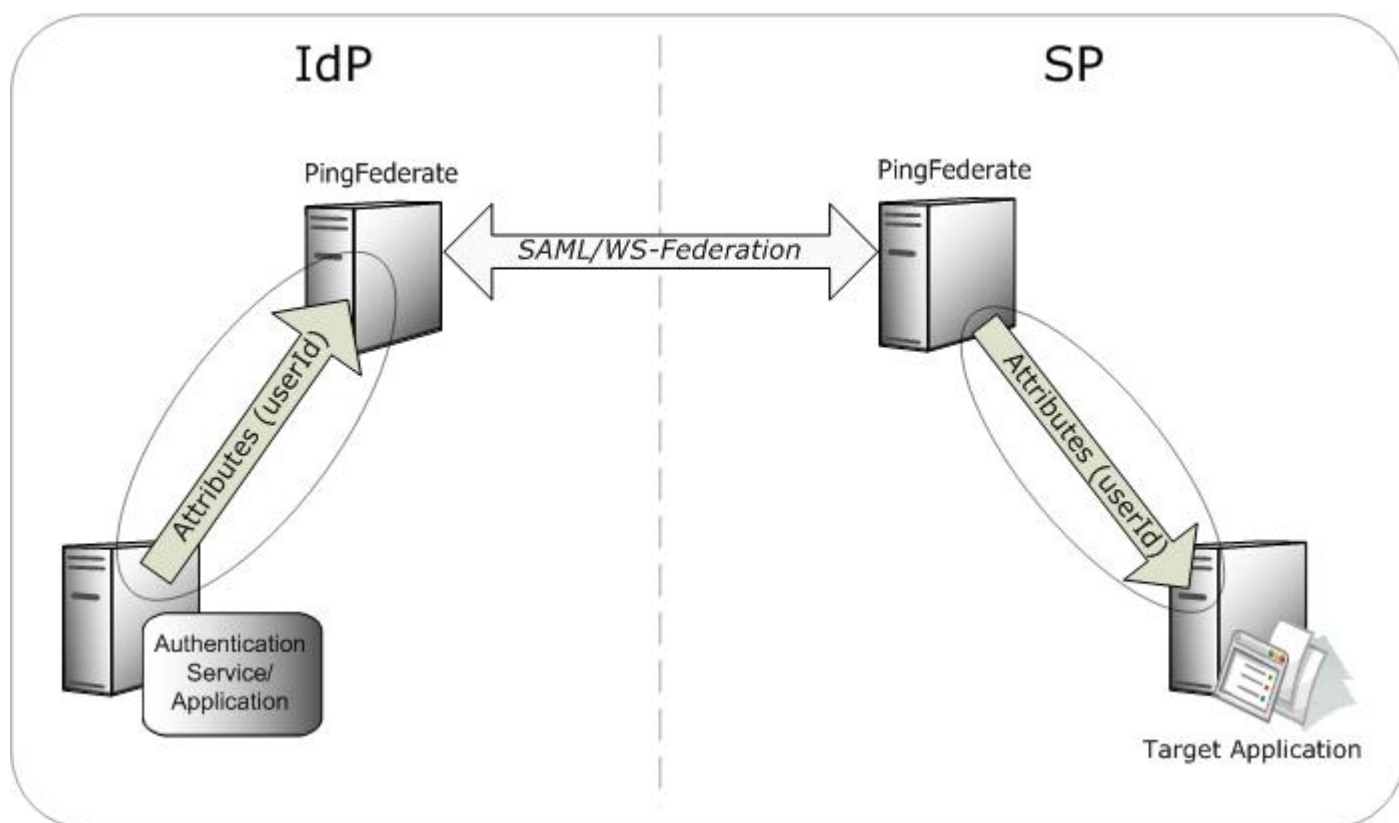
In addition, for integration with the PingFederate WS-Trust security token service (STS), we provide a range of *Token Translators*. These plug-in token processors (for an IdP) and token generators (for an SP) connect the STS with Web Service providers and clients for access to identity-enabled Web Services.

SSO integration concepts

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—Security Assertion Markup Language—refer to [Supported standards](#) on page 29 in the *Get started with PingFederate* guide.) IdP integration typically provides a mechanism through which PingFederate can look up a user's current authenticated session data (for example, 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

An IdP is a system entity that authenticates a user, or “SAML subject,” and transmits referential identity attributes based on that authentication to PingFederate. The IdP integration involves retrieving user-identity attributes from the IdP domain and sending them to the PingFederate server. Typically, the identity attributes are retrieved from an authenticated user session. For IdP integration, a number of attribute-retrieval approaches can be used, depending upon the IdP deployment/implementation environment. Ping Identity offers a broad range of commercial integration kits that address various IdP scenarios, most of which involve either custom-application integration, integration with a commercial IdM product, or integration with an authentication system.

- i **Tip:** For IdPs implementing SSO to selected Software-as-a-Service (SaaS) providers—for example, 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 through application-level integration kits, which 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. The information can be used 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.

Identity management (IdM) systems

An IdP enterprise that uses an IdM system can expand the reach of the IdM domain to external partner applications through integration with PingFederate. IdM integration kits typically use the IdM agent API (if available) 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; integration with PingFederate is accomplished entirely through the PingFederate administrative console.

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

Authentication systems

Initial user authentication is normally handled outside of the PingFederate server using an authentication application or service. PingFederate authentication-system integration kits leverage this local authentication to access applications outside the security domain. For example, an integration kit might access authenticated sessions that are validated against a Windows NTLM or Integrated Windows Authentication (IWA) environment, and then pass session attributes to the PingFederate IdP server.


Authentication integration kits do not require any development; integration with PingFederate is accomplished entirely through the PingFederate administrative console. Ping Identity offers integration kits for authentication systems including:

- IWA/NTLM
- 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 plug-in password credential validators. Supplied validators can use either an LDAP directory, RADIUS server, or a simple username/password verification system maintained by PingFederate. (Customized validators may 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

An SP is the consumer of identity attributes provided by the IdP through a SAML assertion. SP integration involves passing the identity attributes from PingFederate to the target SP application. 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 can involve a number of approaches, and as 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 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. Integration with these custom applications is handled through application-level integration kits, which 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. The information can be used 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 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, ease of deployment and scalability are maximized. Applications running on the

Web/application server must delegate authentication to the server; if the application employs its own authentication mechanism, integration must occur at the application level.

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 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-agent integration kits do not require any development; integration with PingFederate is accomplished entirely through the PingFederate administrative console.

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

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

Identity management (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 (if available) to create an IdM proprietary session token based on the identity attributes received from PingFederate.

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

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

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; integration with PingFederate is accomplished entirely through the PingFederate administrative console.

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

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

Summary

The following table summarizes IdP- and SP-integration deployment scenarios with bundled adapters and some of the integration kits that suit each scenario.

Deployment scenarios	IdP	SP
Custom Application	.NET Integration Kit	.NET Integration Kit

Deployment scenarios	IdP	SP
	Agentless Integration Kit	Agentless Integration Kit
	Java Integration Kit	Java Integration Kit
	PHP Integration Kit	PHP Integration Kit
Identity Management System (IdM)	Web Access Management (WAM) Integration Kit	WAM Integration Kit
Authentication System	Kerberos Adapter (bundled with PingFederate)	Not applicable
	HTML Form Adapter (bundled with PingFederate)	
	HTTP Basic Adapter (bundled with PingFederate)	
	Windows IWA Integration Kit	
	X.509 Integration Kit	
	RSA SecurID Integration Kit	
	Symantec VIP Integration Kit	
Server Agent	NetWeaver Integration Kit	Apache Integration Kit (Red Hat) Apache Integration Kit (Windows) IIS Integration Kit NetWeaver Integration Kit WebSphere Integration Kit
Software as a Service (SaaS)	Not applicable	AWS Connector Box Connector Concur Connector Dropbox Connector Evernote Connector Google Connector Office 365 Connector Salesforce Connector ServiceNow Connector Slack Connector Workday Connector Zendesk Connector

Ping Identity continues to develop new bundled adapters (included with PingFederate) and integration kits; check the Ping Identity [Downloads](#) website for the most up-to-date list of kits.

Please find *User Guides* and other documentation for current integration kits under Ping Identity [documentation](#) website.

SDK developer's guide

The PingFederate® SDK enables integration with IdPs and/or SPs. The interfaces allow developers to build their own custom implementations for communicating authentication and security information between PingFederate and the enterprise environment.

- [Preface](#) on page 559
- [SDK introduction](#) on page 559
- [Get started with the SDK](#) on page 561
- [Implementation guidelines](#) on page 562
- [Build and deploy your project](#) on page 575

Preface

This document provides technical guidance for using the Java Software Development Kit (SDK) for PingFederate. Developers can use this *Guide*, in conjunction with the installed Javadocs, to extend the functionality of the PingFederate server.

Intended Audience

The *Guide* is intended for application developers and system administrators responsible for extending PingFederate, including development of:

- Authentication adapters needed to integrate web applications or identity-management systems (when not already available: see the PingFederate [SSO integration overview on page 552](#), described under [Additional Documentation](#) on page 559.)
- Authentication selectors used to direct SSO authentication to instances of authentication adapters based on specified conditions
- WS-Trust Security Token Service (STS) token translators, including token processors needed to consume and validate security tokens and token generators needed to create security tokens
- Custom data source drivers
- Password credential validators
- Identity store provisioners

The reader should be familiar with Java software-development principles and practices.

Additional Documentation

- Javadocs provide detailed reference information for developers. The Javadocs are located in the `<PF_install>/pingfederate/sdk/doc` directory.
- The PingFederate [SSO integration overview on page 552](#) describes the types of prebuilt authentication adapters Ping Identity provides for integrating web applications and identity-management systems with PingFederate. Since these adapters are based on the SDK, you may want to review this document before building your own adapter to see if your needs have already been met.
- The PingFederate [Administrator's manual on page 57](#) provides background information and user-interface (UI) configuration details needed to integrate implementation(s) of PingFederate interfaces.
- The [user guides for the Java, .NET, and PHP integration kits](#) show examples of SDK implementations.

SDK introduction

The PingFederate Java SDK consists of several application programming interfaces (APIs), including:

- Adapter and STS Token-Translator interfaces
- Authentication selector interfaces
- Custom data source interfaces
- Password Credential Validator interfaces
- Identity Store Provisioner interfaces

Each of these interfaces allows users to create their own plug-ins, customizing certain behaviors of PingFederate to suit an organization's needs. This SDK provides a means to develop, compile, and deploy custom plug-ins to PingFederate.


A number of example plug-ins are included in the PingFederate package for reference. The example projects are located in the `<PF_install>/sdk/plugin-src` directory.

Adapter and STS token-translator interfaces

The adapter and token-translator APIs enable PingFederate integration with IdPs or SPs. The APIs allow developers to build their own custom implementations for communicating authentication and security information between PingFederate and the enterprise environment.

 **Note:** Token-translator interfaces are applicable only to PingFederate versions 6.0 and higher.

In addition to providing requisite runtime integration, an adapter or token translator also describes its configuration parameters to PingFederate; this enables the administrative console to render configuration screens with extensible validation.

 **Note:** Suitable adapter or token-translator implementations for your deployment may already exist, or new implementations may be under development. Before developing your own custom solution, see the Ping Identity [Downloads](#) website for more information about currently 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 connection (for federation hub use cases), depending on the specified conditions. For example, an authentication selector may map internal corporate users to use one adapter, while it maps external non-corporate users to a different adapter.

 **Note:** Authentication selector interfaces are applicable only to PingFederate versions 6.6 and higher.

Authentication selectors are configurable UI plug-ins, allowing you to render custom configuration screens.

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 LDAP or JDBC drivers. This allows developers to retrieve attributes from a data source of their choice during attribute fulfillment for various use cases. Similar to the adapter API, custom data source plug-ins also provide much of the same UI configuration functionality.


Password credential validator interfaces

The password credential validator interfaces allow developers to define credential validators that are used to 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.

 **Note:** Credential validator interfaces are applicable only to PingFederate versions 6.5 and higher.

Identity store provisioner interfaces

Identity Store Provisioners provide a mechanism for provisioning and deprovisioning users to external user stores. For example, a custom Identity Store Provisioner could be configured within an Inbound Provisioning IdP Connection to provision users using the SCIM protocol.

 **Note:** Identity Store Provisioner interfaces are applicable only to PingFederate versions 7.1 and higher.

Similar to the adapter API, Identity Store Provisioners are configurable UI plug-ins, allowing you to render custom configuration screens.

Ping Identity Global Client Services

If you need assistance in using the SDK, visit the Ping Identity [Support](#) website to see how we can help you with your application.

Get started with the SDK

The following sections describe the directories and build components that comprise the SDK and provide instructions for setting up a development environment.

- [Directory structure](#) on page 561
- [Set up your project](#) on page 561

Directory structure

The PingFederate SDK directory (`<PF_install>/pingfederate/sdk`) contains the following:

- `plugin-src/` – The directory where you place your custom plug-in projects. This directory also contains example plug-in implementations showing a wide range of functionality. You may use these examples for developing 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` – This file contains properties used by the Ant build script, `build.xml`, to compile and deploy your custom components. Do not modify this file; use `build.local.properties` to override any properties, if needed.
- `build.local.properties` – Allows you to specify which project you want to build and define properties specific to your environment. The main use of this file is declaring 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.

Set up your project

To start developing your own plug-in:

1. Before you start, ensure you have the Java SDK and Apache Ant installed.
2. To create a new plug-in, create a new project directory in the `<PF_install_dir>/pingfederate/sdk/plugin-src` directory.
3. In the new project directory, create a subdirectory named `java`.

This is where you place the Java source code for your implementation(s).

Follow standard Java package and directory structure layout.

4. If your project depends on third-party libraries, create another subdirectory called `lib` and place the necessary JAR files in it.
5. The build script builds only one project at a time. Edit the `build.local.properties` file and set `target-plugin-name` to specify the name of the directory (under `<PF_install>/pingfederate/sdk/plugin-src`) that contains your project.
6. In `<PF_install>/pingfederate/sdk` run `ant` to display a list of available build targets:

```
[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
[java] jar-plugin Package the plug-in jar
[java]
[java] Default target: help
```

Run the appropriate target to clean, build, or deploy your plug-in.



Note: Building the project with the `build.xml` included in the SDK is recommended because it packages the jars with additional metadata to make it discoverable by PingFederate. For detailed information, see [Build and deploy your project](#) on page 575.

Implementation guidelines

The following sections provide specific programming guidance for developing custom interfaces. Note that the information is not exhaustive—consult the Javadocs to find more details about interfaces discussed here and additional functionality.

Shared interfaces

All plug-in implementations generally invoke methods discussed in the following sections.

- [Configurable plug-in](#) on page 562
- [Describable plug-in](#) on page 562

Configurable plug-in

Any custom plug-in that requires UI settings is considered *configurable* and hence implements the `ConfigurablePlugin` interface. This ensures that PingFederate loads the plug-in instance with the correct configuration settings.

All plug-in types implement the `ConfigurablePlugin` interface and must define the following to enable configuration loading:

```
void configure(Configuration configuration)
```

During processing of a configurable plug-in instance, PingFederate calls the `ConfigurablePlugin.configure()` method and passes in a `Configuration` object. The `Configuration` object provides the plug-in adapter-instance configuration set by an administrator in the PingFederate UI.

The `sp-adapter-example` provided with the SDK shows how to use this method to initialize an adapter-instance from a saved configuration. Once your implementation loads the configuration values, the plug-in instance can use them in other method calls.

Describable plug-in

Any plug-in that requires configuration screens in the PingFederate administrative console is considered a *describable* plug-in. Most plug-ins implement the `DescribablePlugin` interface to ensure that PingFederate renders the correct UI components based on the returned `PluginDescriptor`.

Adapter and custom data source plug-ins are a special case and do not implement the `DescribablePlugin` interface. However, they still return a plug-in descriptor (`AuthnAdapterDescriptor` and `SourceDescriptor` respectively) and are still considered describable plug-ins.

All describable plug-ins must define a UI descriptor. Use one of the following methods to implement a UI descriptor, depending on the type of plug-in:

- For `DescribablePlugin`:

```
PluginDescriptor getPluginDescriptor()
```

- For adapter plug-ins:

```
AuthnAdapterDescriptor getAdapterDescriptor()
```

- For custom data source plug-ins:

```
SourceDescriptor getSourceDescriptor()
```

In many cases, describable plug-ins return a subclass of `PluginDescriptor`, so the return type of the plug-in descriptor getters might be slightly different among plug-in implementations. Your plug-in implementation populates `PluginDescriptor` with `FieldDescriptors`, `FieldValidators`, and `Actions` and is presented as a set of UI components in the PingFederate administrative console.

- **Tip:** Some plug-in types offer concrete descriptor implementations for developers. The Javadocs and examples provided with the SDK show which descriptor classes are available for each plug-in type. The examples also show you how to use `FieldDescriptors`, `FieldValidators`, and `Actions` directly to define your plug-in descriptor.

Implement an IdP adapter

You create an IdP adapter by implementing the `IdpAuthenticationAdapter` or the `IdpAuthenticationAdapterV2` interface. The following Java packages are needed, at a minimum, for implementing this interface:

- `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 interfaces](#) on page 562, you must define the following:

- Session Lookup
- Session Logout

IdP adapter session lookup

```
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
```

PingFederate invokes the `lookupAuthN()` method of your IdP adapter to look up user-session information to handle a request. This method is invoked regardless of whether the request is for IdP- or SP-initiated SSO, an OAuth transaction, or direct IdP-to-SP adapter processing.

- **Note:** The `IdpAuthenticationAdapterV2` interface provides an overloaded version of `lookupAuthN()` applicable to PingFederate versions 6.4 and higher. Use this interface if your adapter requires additional parameters from PingFederate. Refer to the `IdpAuthenticationAdapterV2` interface in the Javadocs for a complete list of available parameters.

In most implementations, a user's session information or a reference to it is communicated to PingFederate via the `HttpServletRequest`, which is passed to the `lookupAuthN()` method. For example, the user's session information can be passed in by 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` in various ways to retrieve the user's session data—for example, by creating a 302 redirect or presenting a web page asking for credentials. If your adapter implementation uses the `HttpServletResponse` to retrieve the user's session information, you must return the user's browser to the URL in the `resumePath` parameter set by the PingFederate runtime server and passed to this method. The `resumePath` is a relative URL signaling PingFederate that a user is continuing an SSO transaction that has already been initiated.

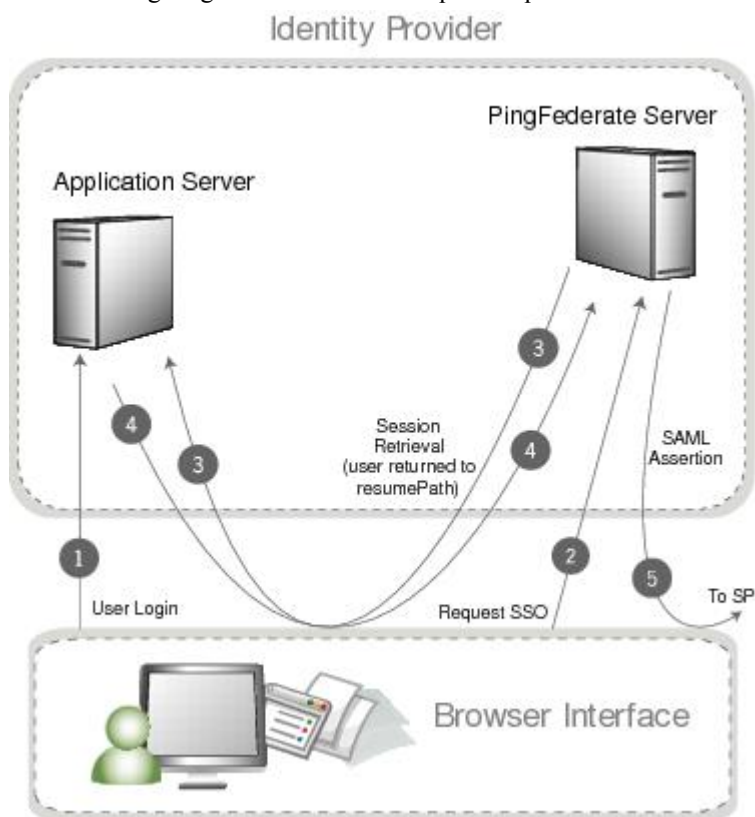
i Tip: When creating a custom adapter, you can design it to render a template for processing and returning HTML to the user's browser using the `TemplateRendererUtil`. A sample (`template-render-adapter-example`) is included in the `sdk/plugin-src` directory of your PingFederate instance.

If your adapter implementation writes to the `HttpServletResponse` to retrieve the user's session data, we recommend that the browser return to the `resumePath` URL at all times, whether the retrieval succeeds or fails. Doing so 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. Such adapters should not be used with the Composite Adapter's “Sufficient” chaining policy (see the “[Composite Adapter Configuration](#)” appendix in the PingFederate *Administrator's manual*).

Processing steps

The following diagram illustrates the request sequence of an IdP-initiated SSO scenario that uses the `resumePath`:



1. User logs in to a local application or domain through an identity-management system or some other authentication mechanism.
2. User clicks a link or otherwise requests access to a protected resource located in the 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. In this example, the adapter determines there is not enough information and redirects the browser to the application server to fetch additional session information.
4. The application server returns the session information and redirects the browser along with the returned information to `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

```
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
```

During 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.

Like the `lookupAuthN()` method, the `logoutAuthN()` method has access to the user's `HttpServletRequest` and `HttpServletResponse` 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. Again, the `resumePath` parameter contains the URL to which the user is redirected to complete the SLO process.

Implement an SP adapter

You create an SP adapter by implementing the `SPAuthenticationAdapter` interface. The Java packages required are, at a minimum:

- `org.sourceid.saml20.adapter.sp.authn`
- `org.sourceid.saml20.adapter.gui`
- `org.sourceid.saml20.adapter.conf`

At a high level, in addition to the methods described under [Shared interfaces](#) on page 562, you must define the following:

- Session Creation
- Session Logout
- Account Linking (if configured in PingFederate for an IdP partner)

SP session creation

```
java.io.Serializable createAuthN(SsoContext ssoContext,
    javax.servlet.http.HttpServletRequest req,
    javax.servlet.http.HttpServletResponse resp,
    java.lang.String resumePath)
```

PingFederate invokes the `createAuthN()` method during the processing of an SSO request to establish a security context in the external application for the user. This method is similar to the `IdpAuthenticationAdapter.lookupAuthN()` method in terms of the objects passed to it and its support for asynchronous requests via the `HttpServletResponse` and `resumePath` parameters. This method also accepts an `SsoContext` object, which has access to information such as user attributes and the target destination URL.

SP adapter session logout

```
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
```

PingFederate invokes the `logoutAuthN()` method during an SLO request to terminate a user's session with the external application. The `HttpServletResponse` and `resumePath` objects are available to support scenarios where redirection of the user's browser is needed to an additional service to clean up any remaining sessions.

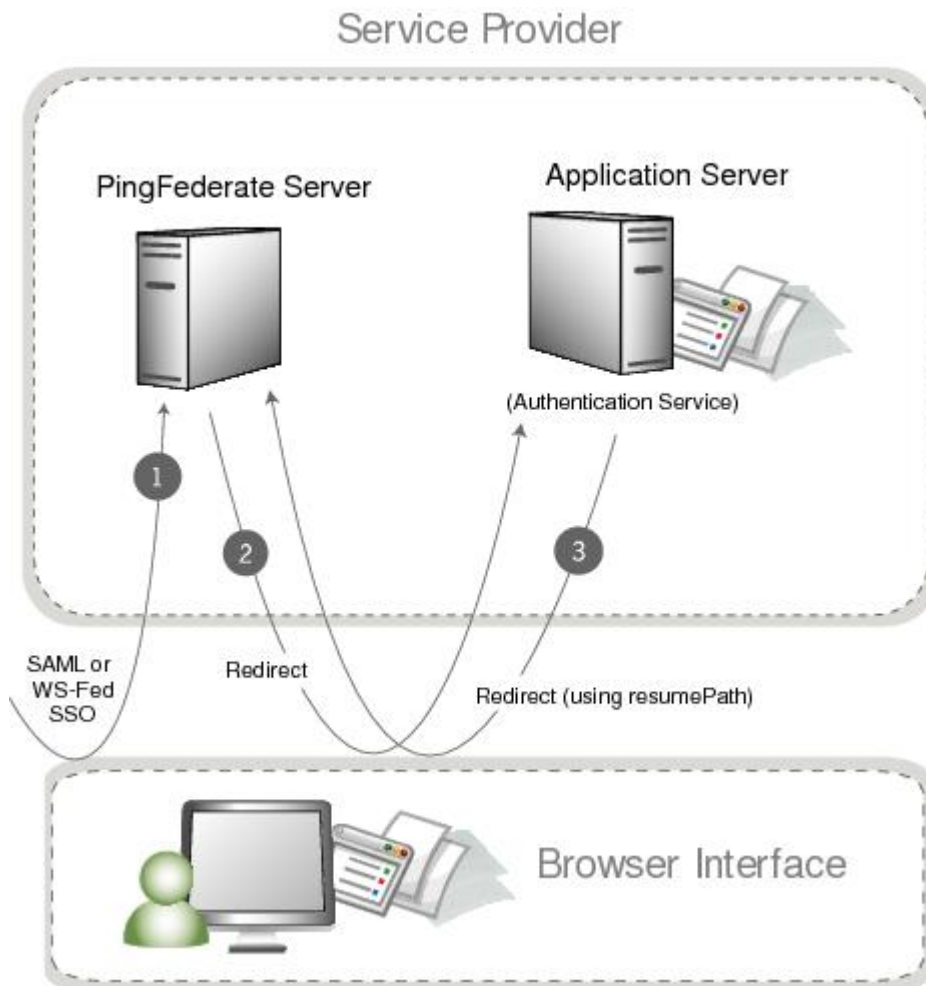
SP account linking

```
java.lang.String lookupLocalUserId(
    javax.servlet.http.HttpServletRequest req,
    javax.servlet.http.HttpServletResponse resp,
    java.lang.String partnerIdpEntityId,
    java.lang.String resumePath)
throws AuthnAdapterException, java.io.IOException
```

PingFederate invokes the `lookupLocalUserId()` method during an SSO request when the IdP connection is configured to use account linking but no account link for this user is yet established. Once the account link is set, PingFederate maintains this information until the user “defederates.” Defederation occurs when the user clicks a link redirecting him/her 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 `HttpServletRequest` to read a local session token. The `String` object returned from the `lookupLocalUserId()` method should be a local user identifier.

Implement a token processor

You create a token-processor implementation (for PingFederate 6.0 and higher) by implementing the `TokenProcessor` interface. The following Java packages are needed, at a minimum, for implementing this 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 interfaces* on page 562, you must define the method:

```
TokenContext processToken(T token)
```

PingFederate invokes the `processToken()` method during the processing of an STS request to perform necessary operations for determining the validity of a token. Type `T` must extend, at a minimum, the type `SecurityToken`. The type `BinarySecurityToken` is also available and may be used to represent custom security tokens that can be transported as Base64-encoded data.

Implement a token generator

You create a token-generator implementation (for PingFederate 6.0 and higher) by implementing the `TokenGenerator` interface. The following Java packages needed, at a minimum, for implementing this 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, described under *Shared interfaces* on page 562, you must define the method:

```
SecurityToken generateToken(TokenContext attributeContext)
```

PingFederate invokes the `generateToken()` method during the processing of an STS request to perform necessary operations for generation of a security token. The type `BinarySecurityToken` is available and may be used to represent custom security tokens that can be transported as Base64-encoded data. The `TokenContext` contains subject data available for insertion into the generated security token.

Implement an authentication selector

Authentication selectors allow PingFederate (version 6.6 and higher) to choose an appropriate authentication source, an 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, the following are the primary Java packages used:

- `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 interfaces* on page 562, you must define the following at a minimum:

- Context Selection
- Authentication Selector Callback

Context selection

```
AuthenticationSelectorContext selectContext(HttpServletRequest req,
    HttpServletResponse resp,
    Map<AuthenticationSourceKey, String> mappedAuthnSourcesNames,
    Map<String, Object> extraParameters,
    String resumePath)
```

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. The `HttpServletRequest` is available to evaluate cookies, parameters, headers, etc. to help determine which authentication source should be selected. The `HttpServletResponse` is also available if the authentication selector requires user interaction to help determine the appropriate authentication source to select. 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 that should be used in conjunction with the `resp` object, such that the user's browser can be sent to this URL to resume the SSO workflow.

Once an authentication source is selected, an `AuthenticationSelectorContext` can be created to denote which authentication source to use. The selected authentication source can be referenced by its ID or by its context. The context is a name that decouples authentication selectors from the configured IDs.

Authentication selector callback

```
void callback(HttpServletRequest req,
    HttpServletResponse resp,
    Map authnIdentifiers,
    AuthenticationSourceKey authenticationSourceKey,
    AuthenticationSelectorContext authnSelectorContext);
```

PingFederate calls the `callback()` method after a selected authentication source is authenticated against. The `callback()` method allows authentication selectors to update resulting attributes, set cookies, or perform other custom functions.



Note: Writing content to the `resp` object in the `callback()` method is not supported, and doing so may result in unexpected behavior. Setting cookies is acceptable.

Implement a custom data source

Out of the box, PingFederate provides the capability of querying data sources for a variety of purposes using LDAP or JDBC interfaces. You can 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 perhaps 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 under *Shared interfaces* on page 562, you must define the following at a minimum:

- Connection Testing

- Available Fields Retrieval
- Data Source Query Handling

Data source connection testing

```
boolean testConnection()
```

When associating a custom data source with an IdP or 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`. Return `false` if your implementation cannot communicate with the data store. 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 could 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. The administrator can then select the attributes from the data source and map 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 that should be populated with data. The method may also receive a `filterConfiguration` object populated with a list of fields. Each field contains a name/value pair that is 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.

You 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.



Note: 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 the 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 that was passed into the method and the values are the attributes.

Implement a password credential validator

Password credential validators allow PingFederate administrators to define a centralized location for username/password validation, allowing validator instances to be referenced by various PingFederate configurations.

To implement a custom password credential validator, the following Java packages need to be imported:

- 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 interfaces* on page 562, you must define the following at a minimum:

```
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, then an `AttributeMap` is returned containing at least one entry representing the principal. If the credentials are invalid, then `null` or an empty map is returned. A `PasswordValidationException` is thrown if the plug-in was unable to validate the credentials (for example, due to an offline host or network problems).

To enable change password in a password credential validator, implement the `com.pingidentity.sdk.password.ChangeablePasswordCredential` interface.


To enable password reset in a password credential validator, implement the `com.pingidentity.sdk.password.ResettablePasswordCredential` interface.

 **Note:** Depending on your password management system, additional system configuration may be necessary to enable password changes—for example, passwords can be changed in Active Directory only if SSL is enabled.

Implement an identity store provisioner


You create an Identity Store Provisioner by implementing the `IdentityStoreProvisionerWithFiltering` or `IdentityStoreProvisioner` interface.

Both interfaces support provisioning and deprovisioning users, and optionally groups, to an external user store. The `IdentityStoreProvisionerWithFiltering` supports list/query and filtering, whereas the `IdentityStoreProvisioner` does not. For more information about list/query and filtering, see [3.2.2 List/Query Resources](#) and [3.2.2.1 Filtering](#) in SCIM Specification (www.simplecloud.info/specs/draft-scim-api-01.html).

 **Note:** As of PingFederate 7.3, the `IdentityStoreUserProvisioner` interface has been deprecated. Developers are encouraged to implement either the `IdentityStoreProvisionerWithFiltering` or `IdentityStoreProvisioner` interface.

Implement the `IdentityStoreProvisionerWithFiltering` interface

Implement the `IdentityStoreProvisionerWithFiltering` interface to provision and deprovision users, and optionally 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.

The following Java packages are needed, at a minimum, for implementing this interface:

- 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.request

 **Note:** Group support is optional (see [Check for group provisioning support](#) on page 572).

For each Identity Store Provisioner implementation, in addition to the methods described under [Shared interfaces](#) on page 562, 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 create-user requests made to `PingFederate` services, for example Inbound Provisioning. This method is responsible for creating the user in the user store managed by the Identity Store Provisioner.

The `CreateUserRequestContext` will contain all information needed to fulfill the request, e.g. user attributes. If the user was successfully provisioned, a `UserResponseContext` should be returned and contain the user attributes used to provision the user. An `IdentityStoreException` should be thrown if an error occurred during the creation process. See `com.pingidentity.sdk.provision.exception` package for different 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, for example Inbound Provisioning. This method is responsible for retrieving user data from the user store managed by the Identity Store Provisioner.

The `ReadUserRequestContext` will contain all information needed to fulfill the request, e.g. user id. If the user data was successfully retrieved, a `UserResponseContext` should be returned and contain the user attributes for the user. 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.

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, for example Inbound Provisioning. This method is responsible for retrieving user data from the user store managed by the Identity Store Provisioner.



Note: The `readUsers` method is applicable only to the `IdentityStoreProvisionerWithFiltering` interface; it does not apply to the `IdentityStoreProvisioner` interface.

The `ReadUsersRequestContext` will contain all information needed to fulfill the request, e.g. filter. If the user data was successfully retrieved, a `UsersResponseContext` should be returned and contain the user attributes satisfying the filter. 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 user

```
UserResponseContext updateUser(UpdateUserRequestContext updateRequestCtx)
throws IdentityStoreException
```

PingFederate invokes the `updateUser()` method of your Identity Store Provisioner in response to update-user requests made to PingFederate services, for example Inbound Provisioning. This method is responsible for updating the user in the user store managed by the Identity Store Provisioner.

The `UpdateUserRequestContext` will contain all information needed to fulfill the request, e.g. user attributes. If the user data was successfully updated, a `UserResponseContext` should be returned containing the user's updated attributes. An `IdentityStoreException` should be thrown if an error occurred during the update process. See `com.pingidentity.sdk.provision.exception` package for different 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 delete-user requests made to PingFederate services, such as Inbound Provisioning. This method is responsible for deprovisioning the user in the user store managed by the Identity Store Provisioner.

The `DeleteUserRequestContext` will contain all information needed to fulfill the request, e.g. user id. An `IdentityStoreException` should be thrown if an error occurred during the deprovision process. See `com.pingidentity.sdk.provision.exception` package for different exceptions that can be thrown.



Note: The plugin implementation for delete MAY 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`.

Check for group provisioning support

```
boolean isGroupProvisioningSupported()
throws IdentityStoreException
```

Implement this `isGroupProvisioningSupported()` method to return true if group provisioning is supported by your Identity Store Provisioner or false otherwise. An `IdentityStoreException` should be thrown if an error occurred during the query process. See `com.pingidentity.sdk.provision.exception` package for different exceptions that can be thrown.

Create group

```
GroupResponseContext createGroup(CreateGroupRequestContext
createRequestCtx)
throws IdentityStoreException
```

PingFederate invokes the `createGroup()` method of your Identity Store Provisioner in response to create-group requests made to PingFederate services, for example Inbound Provisioning. This

method is responsible for creating the group in the user store managed by the Identity Store Provisioner if the `isGroupProvisioningSupported()` returns true; otherwise, it should throw `NotImplementedException`.

The `CreateGroupRequestContext` will contain all information needed to fulfill the request, e.g. the group attributes. If the group was successfully provisioned, a `GroupResponseContext` should be returned and contain the group attributes used to provision the group. An `IdentityStoreException` should be thrown if an error occurred during the creation process. See `com.pingidentity.sdk.provision.exception` package for different exceptions that can be thrown.

Read group

```
GroupResponseContext readGroup(ReadGroupRequestContext readRequestCtx)
throws IdentityStoreException
```

`PingFederate` invokes the `readGroup()` method of your Identity Store Provisioner in response to get-group requests made to `PingFederate` services, for example Inbound Provisioning. This method is responsible for retrieving group data from the user store managed by the Identity Store Provisioner if the `isGroupProvisioningSupported()` returns true; otherwise, it should throw `NotImplementedException`.

The `ReadGroupRequestContext` will contain all information needed to fulfill the request, e.g. group id. If the group data was successfully retrieved, a `GroupResponseContext` should be returned and contain the group attributes for the group. 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.

Read groups

```
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, for example Inbound Provisioning. This method is responsible for retrieving 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 is applicable only to the `IdentityStoreProvisionerWithFiltering` interface; it does not apply to the `IdentityStoreProvisioner` interface.

The `ReadGroupsRequestContext` will contain all information needed to fulfill the request, e.g. 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, for example Inbound Provisioning. This method is responsible for updating the group in the user store managed by the Identity Store

Provisioner if the `isGroupProvisioningSupported()` returns true; otherwise, it should throw `NotImplementedException`.

The `UpdateGroupRequestContext` will contain all information needed to fulfill the request, e.g. group attributes. If the group data was successfully updated, a `GroupResponseContext` should be returned containing the group's updated attributes. An `IdentityStoreException` should be thrown if an error occurred during the update process. See `com.pingidentity.sdk.provision.exception` package for different 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 is responsible for deprovisioning the group in the user store managed by the Identity Store Provisioner if the `isGroupProvisioningSupported()` returns true; otherwise, it should throw `NotImplementedException`.

The `DeleteGroupRequestContext` will contain all information needed to fulfill the request, e.g. group id. An `IdentityStoreException` should be thrown if an error occurred during the deprovision process. See `com.pingidentity.sdk.provision.exception` package for different exceptions that can be thrown.

Implement the IdentityStoreUserProvisioner interface



Note: The `IdentityStoreUserProvisioner` interface has been deprecated since `PingFederate 7.3`. Developers are encouraged to implement either the `IdentityStoreProvisionerWithFiltering` or `IdentityStoreProvisioner` interface.

Implement the `IdentityStoreUserProvisioner` interface to provision and deprovision users to an external user store.



Tip: The `IdentityStoreUserProvisioner` interface does not provision or deprovision groups. For group support, see [Implement the IdentityStoreProvisionerWithFiltering interface](#) on page 570.

The following Java packages are needed, at a minimum, for implementing this 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 interfaces](#) on page 562, you must implement the following:

- 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 create-user requests made to `PingFederate` services, for example Inbound Provisioning. This method is responsible for creating the user in the user store managed by the Identity Store Provisioner.

The `CreateUserRequestContext` will contain all information needed to fulfill the request, e.g. user attributes. If the user was successfully provisioned, a `UserResponseContext` should be returned and contain the user attributes used to provision the user. An `IdentityStoreException` should be thrown if an error occurred during the creation process. See `com.pingidentity.sdk.provision.exception` package for different 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, for example Inbound Provisioning. This method is responsible for retrieving user data from the user store managed by the Identity Store Provisioner.

The `ReadUserRequestContext` will contain all information needed to fulfill the request, e.g. user id. If the user data was successfully retrieved, a `UserResponseContext` should be returned and contain the user attributes for the user. 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 user

```
UserResponseContext updateUser(UpdateUserRequestContext updateRequestCtx)
throws IdentityStoreException
```

`PingFederate` invokes the `updateUser()` method of your Identity Store Provisioner in response to update-user requests made to `PingFederate` services, for example Inbound Provisioning. This method is responsible for updating the user in the user store managed by the Identity Store Provisioner.

The `UpdateUserRequestContext` will contain all information needed to fulfill the request, e.g. user attributes. If the user data was successfully updated, a `UserResponseContext` should be returned containing the user's updated attributes. An `IdentityStoreException` should be thrown if an error occurred during the update process. See `com.pingidentity.sdk.provision.exception` package for different 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 delete-user requests made to `PingFederate` services, such as Inbound Provisioning. This method is responsible for deprovisioning the user in the user store managed by the Identity Store Provisioner.

The `DeleteUserRequestContext` will contain all information needed to fulfill the request, e.g. user id. An `IdentityStoreException` should be thrown if an error occurred during the deprovision process. See `com.pingidentity.sdk.provision.exception` package for different exceptions that can be thrown.



Note: The plugin implementation for delete MAY 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`.

Build and deploy your project

To build and deploy your project, you can choose to use the provided Apache Ant script or another build utility.

Build and deploy with Ant

The PingFederate Java SDK comes with an Apache Ant build script that makes building and deploying your project simple.

1. Edit the `build.local.properties` file and set the `target-plugin.name` property to the name of your project subdirectory (see [Directory structure](#) on page 561).



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 jars, place them into your project's `lib` directory.

If the directory does not exist, create a new directory called `lib`, 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.

To clean the project, enter:

```
ant clean-plugin
```

To compile the project, enter:

```
ant compile-plugin
```

To compile the project and create a JAR, enter:

```
ant jar-plugin
```

The SDK creates deployment descriptor(s) in the `PF-INF` directory and places it in a JAR. The descriptor tells PingFederate what plug-in implementations are contained in the JAR.

The compiled class files and the deployment descriptor(s) are placed in the `pingfederate/sdk/plugin-src/<subproject-name>/build/classes` directory.

The `pf.plugins.<subproject-name>.jar` file is placed in the `pingfederate/sdk/plugin-src/<subproject-name>/build/jar` directory.

To compile, create a JAR, and deploy the project to PingFederate, enter:

```
ant deploy-plugin
```

This build target performs the steps described above and deploying any JAR files found in the `lib` directory of your subproject.



Note: To deploy your plug-in 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.

Build and deploy manually

To build your project with another build utility, you must take some prerequisite steps to create the deployment descriptors for each of your plug-ins. The deployment descriptor files allow PingFederate to discover your plug-ins.

Create deployment descriptors

1. In your project, create a new directory called `PF-INF`. This directory must be at the root of your JAR file, similar to `META-INF`.
2. Inside `PF-INF` create the appropriate text file(s) for each type of plug-ins you created:

Plug-in Type	Filename
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
Identity Store Provisioner	identity-store-provisioners

- In each text file created, specify the fully qualified class name of each plug-in that implements the corresponding plug-in interface. Place each class name on a separate line.

Build your project manually

- To compile your project, you need to have the following directories on your classpath:
 - `pingfederate/server/default/lib`
 - `pingfederate/lib`
 - `pingfederate/sdk/lib`
 - `pingfederate/sdk/plugin-src/<subproject-name>/lib`
- To create a JAR, simply archive the compiled class files along with the deployment descriptor(s) using your build tool. The deployment descriptors must be in the `PF-INF` directory, located at the root of the JAR.

Deploy your project

- To deploy your plug-in, simply copy the JAR file and any third-party JAR files into the `<pf_install>/pingfederate/server/default/deploy` directory of the PingFederate installation.

Logging

You can 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 SP adapter contained in the directory `sdk/adapters-src/sp-adapter-example` shows how to use a logger for your adapter.

Release notes

PingFederate 8.4.4 - November 2017

PingFederate® 8.4.4 is a cumulative maintenance release for PingFederate 8.4, which introduced many new features, such as configuration at scale, OpenID Connect signed requests, JWT OAuth 2.0 authorization grants and client authentication, dynamic discovery of cluster nodes using AWS Roles, self-service account unlock, Docker support, and more. For a summary of the 8.4 release, see the [PingFederate 8.4 - June 2017](#) on page 588.

Resolved issues

Ticket ID	Description
PF-17472	Fixed an issue where the administrative API failed to update an IdP connection when it was configured with account linking and OAuth attribute mapping.
PF-17349 and PF-17330	Resolved two rare deadlock issues.
PF-17215	Resolved an issue where a configuration replication error could occur when a previously removed OAuth client was recreated with the same client ID.
PF-17127	Introduced a new setting (<code>InterReqStateMgmtMapImpl.max.session.attrs</code> in the <code>size-limit.conf</code> file) to optimize the memory usage of the Inter-Request State-Management Service.
PF-17108	Resolved an issue where the Kerberos Adapter could not fail over to the next adapter when the <code>pf.runtime.context.path</code> property was configured in the <code>run.properties</code> file.
PF-17035	Fixed an issue where the Composite Adapter could skip the OpenToken Adapter's Logout Service URL when processing logout requests.
PF-17009	Resolved an issue where expired grants were not cleaned up when they were stored on an LDAP server.

Upgrade considerations

Several specific modifications made since PingFederate® 8.0 may affect existing deployments, as described in the following sections.

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 [Copy customized files or settings](#) on page 630 for more information.

An improved index (`IDX_FIELD_NAME`) in the database table for OAuth clients

PingFederate 8.4 modified an existing index (`IDX_FIELD_NAME`) in the `pingfederate_oauth_clients_ext` database table as a general improvement. For information on modifying this index in your existing table, see [Review database changes](#) on page 631.

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`) via 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 could be re-enabled. Refer to [this knowledge base article](#) 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, to which is referred 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 [Configure cluster protocol properties](#) on page 535.

Metadata signing

Previously, when no signing certificate was chosen in the **Server Configuration > Server Settings > Metadata Signing** screen, 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 in the **Server Configuration > Signing & Decryption Keys & Certificates** screen.

Starting with PingFederate 8.1, if no certificate is selected in the **Metadata Signing** screen, PingFederate provides unsigned metadata at both aforementioned endpoints. Select a certificate in the **Metadata Signing** screen 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 for greater 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 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 4.0 (or older) drivers, 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 [Review database changes](#) on page 631.

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 [Review log configuration](#) on page 633 for instructions.



Note: PingFederate has been tested with vendor-specific JDBC 4.1 drivers. To obtain your database driver JAR file, contact your database vendor. Database driver file should be installed to the `<pf_install>/pingfederate/server/default/lib` directory. You must restart the server after installing the driver.

Upgrade considerations introduced in PingFederate 7.x

Hostname verification for LDAPS

For LDAP type data stores with LDAPS enabled, hostname verification of the certificate is enabled by default for all new data stores starting with PingFederate 7.3. When upgrading from a previous version of PingFederate, this option is disabled for existing data stores 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 data store. Follow the instructions in [Review database changes](#) on page 631 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 multi-valued attributes in WS-Federation assertions

Starting with PingFederate 7.3, multi-valued 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 [Review database changes](#) on page 631.

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 [Review database changes](#) on page 631).

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. For more information, see [HTML Form Adapter](#) on page 430.

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 [Review database changes](#) on page 631.

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.

Multi-valued 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 multi-valued 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 may be reverted via the `supportMultiValuesFromDirectory` property located in the `<pf_install>/pingfederate/server/default/data/config-store/com.pingidentity.provisioner.mapping.OgnlFieldMapper.xml` file.

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 may 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 may 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.

 **Note:** Although strongly discouraged, this policy change may 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.HandleAuthnRequest.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 such connections to use the newer algorithm RSA-OAEP (see [Select an encryption certificate \(SAML 2.0\)](#) and [Choose an encryption certificate \(SAML 2.0\)](#)).

 **Note:** The JCE provider libraries distributed with SafeNet Luna SA HSM does not support the RSA-OAEP algorithm at the time of the PingFederate 7.1 R2 release. As a result, RSA-v1.5 is still allowed when deploying PingFederate with a Luna SA HSM.

OAuth persistent grants expiration

When upgrading to PingFederate 6.8 or higher, 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.

Deprecated features

Logging configuration

The default logging configuration has been vastly optimized since PingFederate® 8.2. As a result, the product distribution no longer includes the `terse.example.log4j2.xml` file in the `<pf_install>/pingfederate/server/default/conf` directory, starting with version 8.4.

For more information, see [Log4j 2 logging service and configuration](#) on page 96.

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 (if any), you must migrate custom changes manually. For more information, see [Copy customized files or settings](#) on page 630.

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 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 4.0 (or older) drivers, contact your vendors for the latest drivers and replace the older JDBC database-driver JAR files with the latest.

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 (WS-Trust STS use cases) 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.

For instructions, see [Migrate to the integrated Username Token Processor](#) on page 634.

JMX monitoring support for outbound provisioning

As of PingFederate 7.2, the JMX monitoring support for outbound provisioning has been deprecated and replaced with an audit file approach that is more inline with other auditing services offered by PingFederate.

JMX monitoring for outbound provisioning can be re-enabled in PingFederate 7.2 but will be removed in a future release (see [Monitor outbound provisioning](#) on page 149 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.

Known issues

Self-service password reset

- Passwords can be reset for AD accounts without the permission to change password.

Administrative console and administrative API

- With respect to managing IdP or SP connections, the administrative API connection support is currently limited to Browser SSO and SLO connections. As a result, it is possible to lose settings that are not supported by the administrative API (for example, WS-Trust and provisioning) when these settings are added through the administrative console (the UI) and the connection is subsequently modified through the API.

Known limitations

Administrative console and administrative API

- Prior to toggling the status of a connection via 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 plug-in, 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 screens in the PingFederate console.
- If authenticated to the PingFederate administrative console using certificate authentication, a session that has timed out may 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 to the Main Menu once authentication is complete. Similar behavior applies for certificate authentication, in principle. However, because the browser may automatically resubmit the certificate for authentication, the browser may redirect to the Main Menu and not the login page.

Hardware security modules

- When PingFederate is integrated with SafeNet Luna SA on a platform with Oracle Java SE Runtime Environment (Server JRE) 8 update 102, runtime errors may occur when handling certificates with a signing algorithm of RSA SHA256, SHA384, or SHA512. After switching to Oracle Server JRE 8 update 91, the issue is no longer observed.
- When using PingFederate with the Thales nShield Connect HSM, it is not possible to use an elliptic curve (EC) certificate as an SSL server certificate.

Composite Adapter configuration

- SLO is not supported when users are authenticated through a Composite Adapter instance that contains another instance of the Composite Adapter.
- Adapter instances specified as **Sufficient** in a Composite Adapter configuration should be limited to adapter types that explicitly return control to PingFederate after a failure. Otherwise, the next adapter instance in the authentication "chaining" sequence (if any) may not be tried, and other unexpected behavior may occur. The following adapters work correctly under the **Sufficient** authentication policy in failure mode:
 - HTML Form
 - HTTP Basic
 - Kerberos
 - CoreBlox
 - Duo Security
 - IWA (returns control to PingFederate only if the failure is the result of invalid credentials after the configured number of retries)
 - OpenID (Generic)
 - OpenID (Google)
 - X.509

Note that this list is updated as other adapters are modified, tested, and released.

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.

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.

Provisioning

- When using a filter to query the `/pf-scim/v1/Users` or `/pf-scim/v1/Groups` endpoints for user or group information, if the user repository is Microsoft Active Directory 2008 R2, sorting in descending order does not work. Without the filter, the results can be sorted in descending order. Consider the following examples:

`https://www.example.com/pf-scim/v1/Users?filter=userName%20sw%20%22j%22&sortBy=userName&sortOrder=descending` does not sort the results in descending order.

`https://www.example.com/pf-scim/v1/Users?filter=userName%20sw%20%22j%22&sortBy=userName&sortOrder=ascending` sorts the results in ascending order, as expected.

<https://www.example.com/pf-scim/v1/Users?sortBy=username&sortOrder=descending> sorts the results in descending order, as expected.

Note that this limitation does not apply when the user repository is Microsoft Active Directory 2012.

- 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.

Database logging

- If PingFederate cannot establish a 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.

Note that 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.

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.

Configcopy

- If you are using `configcopy` to copy all connections, channels, data sources, adapters, or token translators and you choose to set override properties, the override is applied to all instances. It is recommended that you use care when applying overrides for copy-all operations.
- The `configcopy` tool supports copying only a single reference for each of the following that are defined for a given connection: adapter, data source, Assertion Consumer Service URL, Single Logout Service URL, and Artifact Resolution Service URL. If you have multiple adapters, data stores, or any of the aforementioned service URLs 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 you choose to set an override parameter 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 plug-in configurations (including adapters, token translators, and custom data stores), does not currently support overrides of complex data structures, including tables, extended contract attributes, and masked fields.
- When using `configcopy` to copy connection data, any SOAP 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.

Previous releases

PingFederate 8.4.3 - October 2017

PingFederate® 8.4.3 is a cumulative maintenance release for PingFederate 8.4, which introduced many new features, such as configuration at scale, OpenID Connect signed requests, JWT OAuth 2.0 authorization grants and client authentication, dynamic discovery of cluster nodes using AWS Roles, self-service account unlock, Docker support, and more. For a summary of the 8.4 release, see the [PingFederate 8.4 - June 2017](#) on page 588.

Resolved issues

Ticket ID	Description
PF-17097	Resolved an issue where the outbound provisioning engine could stop and subsequently wrote a "Timer already cancelled" error message to the provisioner log.

Ticket ID	Description
PF-17053	Fixed an issue where SSO requests might fail when connections were configured with one or more ReferenceID Adapter instances with per-connection overridden instance settings.
PF-16997	If an SSO request invokes a target URL-to-SP adapter instance entry (configured on the SP Configuration > Target URL Mapping screen), and if the resulting SP adapter instance is one of the multiple target session mappings with per-connection overridden instance settings, PingFederate now honors the overridden values at runtime.

PingFederate 8.4.2 - September 2017

PingFederate® 8.4.2 is a cumulative maintenance release for PingFederate 8.4, which introduced many new features, such as configuration at scale, OpenID Connect signed requests, JWT OAuth 2.0 authorization grants and client authentication, dynamic discovery of cluster nodes using AWS Roles, self-service account unlock, Docker support, and more. For a summary of the 8.4 release, see the [PingFederate 8.4 - June 2017](#) on page 588.

Resolved issues

Ticket ID	Description
PF-16908	Fixed an issue where the Kerberos Adapter returned the Username attribute value as that of the Domain/Realm Name attribute.
PF-16878	Resolved an issue where some OAuth clients that had been granted access to an Access Token Management instance in the past might lose such privilege after an upgrade.
PF-16876	Upgraded to Jetty 9.3.20 to resolve an issue that adversely affected the daily log rotation process for the Jetty request log files.
PF-16733	Lessen the lifetime of the inter-request state information to match its short-lived nature and thus reduced the memory footprint of PingFederate. (The setting is stored in the <code>size-limits.conf</code> file. For more information, see Upgrade considerations on page 578.)
PF-16683	Fixed an issue where users were not able to abort their requests for self-service password reset or continue with their requests under certain circumstances.
PF-16559	Resolved a rare upgrade issue where WS-Trust STS connections could be affected adversely.
PF-16514	Added a per-connection option to disregard a validation condition where 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. For more information, please refer to the <code>org.sourceid.saml20.util.VirtualIdentityUtil.xml</code> file.
PF-16475	Resolved an issue where configuration errors might occur to connections that are created with identical Partner's Entity ID value when letter case is ignore; for example, EXAMPLE.COM versus example.com.
PF-16455 and PF-16412	PingFederate is now more resilient when loading erroneous connections after an upgrade or during an import of a configuration archive. For instance, if multiple connections of the same role are using the same entity ID, PingFederate retains the latest connection automatically. Additionally, if PingFederate encounters incomplete connections, it ignores such entries without halting the rest of

Ticket ID	Description
PF-16436	<p>the process. Furthermore, such rare error conditions are now recorded in the server log, giving administrators the opportunity to review them at a later time.</p> <p>Resolved a user interface issue in the administrative console where the console and the administrative API were configured to use native authentication and certificate authentication, respectively.</p>

PingFederate 8.4.1 - August 2017

PingFederate® 8.4.1 is a cumulative maintenance release for PingFederate 8.4, which introduced many new features, such as configuration at scale, OpenID Connect signed requests, JWT OAuth 2.0 authorization grants and client authentication, dynamic discovery of cluster nodes using AWS Roles, self-service account unlock, Docker support, and more. For a summary of the 8.4 release, see the [PingFederate 8.4 - June 2017](#) on page 588.

Resolved issues

Ticket ID	Description
PF-16364	Resolved an issue where requests would fail if they invoked an authentication policy path that contained different instances of the Reference ID Adapter as back-to-back checkpoints.
PF-16329	The optional authentication sessions feature has been enhanced to handle scenarios where state servers may not always contain the same session-state information.
PF-16260	Resolved an issue where custom implementation of the OAuth Client Storage plug-in could fail after an upgrade.
PF-16190	When OAuth 2.0 Bearer Token is the chosen authentication scheme for SCIM outbound provisioning, PingFederate now transmits request parameters in the message body.
PF-16141	The OpenID Provider configuration endpoint (<code>/well-known/openid-configuration</code>) now returns the <code>request_object_signing_alg_values_supported</code> value in the format of a regular string array.
PF-16119	Resolved a thread-safety issue in the Kerberos Adapter.
PF-16069	Improved the administrative console experience. Administrators are no longer required to clear the browser cache files in order for their browsers to render the user interface properly.
PF-15813	Resolved an issue where the administrative console did not allow an administrator to save an OAuth Assertion Grant IdP connection when WS-Trust was the only SP protocol that had been enabled (on the Server Configuration > Server Settings > Roles & Protocols screen).
PF-15760	Upgraded to jose4j 0.5.7.
PF-15709	The SCIM inbound provisioning endpoints (<code>/pf-scim/v1/Users</code> and <code>/pf-scim/v1/Groups</code>) can now handle responses with over 10,000 records.
PF-15556	Resolved an issue where users of Internet Explorer were not able to fall through an instance of the Kerberos Adapter (or an instance of the IWA IdP Adapter that had been configured to allow Kerberos authentication only) to the next checkpoint in an authentication policy.

PingFederate 8.4 - June 2017

Enhancements

Configuration at scale

We have set a new bar for the number of IdP connections, SP connections, and OAuth clients that our customers can configure. We looked at the number of connections and OAuth clients that some of our largest customers were using and made improvements so that their investment could grow at least tenfold. This enhancement includes updates to paging, sorting, filtering, and searching in the administrative console to make it easier to find connections and OAuth clients. Configuration replication is also optimized to only replicate the connections that have changed, greatly reducing the amount of data that needs to be distributed and processed throughout a cluster when configuration changes.

Additionally, a new option is available to schedule the creation of backup configuration archives as an alternative to the generation of archives at sign-on to the administrative console. This option improves the responsiveness of the administrative-console sign-on experience for customers with many connections and OAuth clients because the backup process could take more than a few seconds.

The SDK support for custom storage of OAuth clients has been extended with a new interface, `ClientStorageManagerV2`. This interface includes a `search()` method, allowing customers to provide efficient implementations of the pagination and search functions exposed in the administrative console.

OpenID Connect signed requests

OpenID Connect defines a method to sign authentication requests sent from Relying Parties (RPs) so that the OpenID Provider (OP) can ensure that the request has not been tampered with in transit (see https://openid.net/specs/openid-connect-core-1_0.html#JWTRequests). We've added support for this in PingFederate, and this will open up new use cases. With signed requests, RPs can convey contextual information about the specific transaction the user is attempting. Authorization servers can then use that information to drive different interactions with the user, such as prompting the user for fine-grained consent. Customers can either implement a custom adapter to gain access to the data in the request parameter or look forward to future adapters shipped by Ping Identity that will expose this data. Further, we added support for both requiring signed requests from RPs and sending signed requests to OPs.

JWT OAuth 2.0 authorization grants

We have implemented support of the JWT profile for OAuth 2.0 authorization grants (see <https://tools.ietf.org/html/rfc7523#section-2.1>). The primary use case for this grant type is for cross-domain API transactions where the transaction is within the context of a user's consent but the user is not directly present. To use this grant type an OAuth client submits a signed JWT to an authorization server and receives an access token. This feature aligns with our existing support for the SAML assertion profile for OAuth 2.0 authorization grants.

To further support this use case we added a related feature to help customers that have OAuth clients that make cross-domain API calls and need to use this grant type with a remote AS. Clients that make calls to partner APIs that need to use a JWT authorization grant can exchange a local token for the necessary JWT using WS-Trust. This again aligns with our support of the equivalent SAML assertion profile.

JWT OAuth 2.0 client authentication

Another OAuth profile that is related to the JWT authorization grant, which is defined in the same specification, is OAuth 2.0 client authentication (see <https://tools.ietf.org/html/rfc7523#section-2.2>). We have also added support for this, and it allows OAuth clients to authenticate to an authorization server using a signed JWT. This provides a better security model than client ID and secret authentication as the private key used to sign the JWT remains secret to the client and does not need to be shared with an outside system. This support has been added for when PingFederate is the authorization server or when PingFederate is the OpenID Connect Relying Party.

Dynamic discovery of cluster nodes using AWS Roles

We've expanded our support for dynamic discovery (to support elastic scaling) of cluster nodes in AWS by leveraging the permissions of an AWS Role that's been assigned to an AMI. You can now use this option as an alternative to the existing option of setting an AWS IAM access key and secret for dynamic discovery. This new

feature also adds the ability for PingFederate to discover nodes based on static and dynamic AWS AMI metadata through configuration of tags and filters.

Self-service account unlock

Adding to the existing self-service password reset feature of the HTML Form Adapter, we introduce a new feature to allow users to unlock accounts that have been locked due to password policy at the underlying LDAP directory. This optional feature can help in situations when users realize that they have mistyped their passwords several times and simply want to unlock their accounts, preserving their existing passwords.

Cluster-aware account lockout protection

Account lockout protection prevents user accounts from becoming locked at the underlying user repository based on too many failed authentication attempts. This protection, which is available in many areas of PingFederate® (for example, the HTML Form Adapter, the Username Token Processor, the OAuth Resource Owner Password Credentials grant type, and the administrative console native authentication scheme), is now shared across nodes in a cluster. This enhancement helps in situations where PingFederate is deployed behind a load balancing infrastructure without sticky sessions.

Message customization and localization for password reset and account unlock

Administrators now have new options for customizing and localizing messages sent to users for self-service password reset and account unlock. Similar to the existing support that we already have for customizing email messages, SMS messages can also be customized. Furthermore, these messages can be localized based on the user's preferred language, captured when the user initiates self-service password reset or account unlock.

Docker support

We continue to add features to help customers automate the deployment of PingFederate. With this release we now officially qualify on Docker.

Configurable cluster encryption strength

Cluster encryption strength is now configurable in PingFederate. The encryption strength varies depending on the cryptographic providers available from the Oracle Java SE Runtime Environment (Server JRE).


Security enhancement

The default TLS cipher suites for HTTPS listeners have been updated to follow the latest security best practices.

Other improvements

- Integration with Microsoft Office 365 has been improved to better support Microsoft Intune and Dynamics 365.
- PostgreSQL relational database support has been added and can be used for attribute lookup, OAuth persistent grant storage, OAuth client storage, logging, account linking, and outbound provisioning.
- Improved user provisioning efficiency to reduce execution time when multiple channels are configured.
- Tracking IDs are now included by default in audit log entries.
- Header based certificate authentication has been improved to be tolerant of missing PEM headers.
- The following bundled components and third-party dependencies have been updated:
 - Log4j 2 2.8.2
 - Jackson 2.7.8
 - Jetty 9.3.16
 - jose4j 0.5.5
 - Twilio Java SDK 7.4
 - UnboundID LDAP SDK for Java 3.2

Resolved issues

Ticket ID	Description
PF-15778	Fixed an access token validation issue where any claim with a data type other than string or list was not structured properly in the JSON response.
PF-15757	Resolved an issue where PingFederate did not update or clear the internal data store after changes had been made in the attribute mapping configuration for outbound provisioning.
PF-15756	The Kerberos Adapter now uses the correct Velocity template when a Kerberos Adapter instance is used as one of the checkpoints in an authentication policy and the authentication attempt has failed.
PF-15755	Authentication policies can now complete the failover process when an IdP adapter renders a Velocity template while returning a failure response status.
PF-15708	EXCEPTION messages from the <code>com.pingidentity.fsm.state.impl.PortalMenuState.getClientCountError</code> class are now categorized as ERROR (log level) in the server log.
PF-15630	Fixed an issue where the self-service password reset feature did not use information from the <code>pf-accept-language</code> cookie (if presented) to identify the locale of the user.
PF-15629	Resolved an OAuth client configuration issue where changes made to client authentication scheme were not saved occasionally.
PF-15564	Addressed an issue where contextual attribute values were not passed to the follow-on IdP connections when SP Authentication Policies were enabled
PF-15327	The <code>kerberos.error.template.html</code> template file now supports localization.
PF-15286	The <code>html.form.login.template.headerMessage</code> entry has been removed from the <code>pingfederate-message.properties</code> file as it is no longer applicable.
PF-15258	Resolved an issue where the username field values were always masked in the audit log for SLO requests through any OpenID Connect IdP connections.
PF-15201	Improved the validation process for the Hostname(s) field value in the LDAP data store configuration screen.
PF-15009	Improved the validation process for signed SAML 2.0 AuthnRequest messages; the AssertionConsumerServiceURL element value must be a full URL.
PF-14946	When processing SAML response messages that do not contain the Subject element, PingFederate now handles this error condition accordingly.
PF-14876	<p>Updated the <code>log4j2.xml</code> configuration file with a new logger (<code>com.pingidentity.common.util.xml.XmlBeansUtil</code>).</p> <p> Caution: This new logger is disabled by default. When enabled, PingFederate may write sensitive user information to the server log. Consider enabling this logger for the sole purpose of troubleshooting in non-production environments only (and disabling this logger when it is no longer required).</p>
PF-14602	Error conditions between the OAuth <code>/as/token.oauth2</code> endpoint and the clients are now better captured in the audit and server logs.
PF-14545 and PF-13334	When both signature verification and encryption are disabled, PingFederate now removes such partner certificates from the connection.
PF-14411	Fixed a dependency error when at least one attribute in an authentication policy contract deployed in an authentication policy was fulfilled by an OGNL expression.

Ticket ID	Description
PF-14361	Administrators must now select at least one virtual server ID when the Restrict Virtual Server IDs check box is selected on the Virtual Server IDs screen.
PF-14041	Improved the outbound provisioning process to handle a scenario where the Timestamp Attribute of an object does not contain any value.
PF-13733	Account lockout protection (via the Account Locking Service) has been extended to the Username Token Processor (UTP). PingFederate now rejects authentication attempts for a configurable amount of time after three failures. As needed, administrators may also override the failure threshold value per UTP instance.
PF-13698	When updating a SAML 2.0 connection by metadata, if encryption has not been enabled prior to the update, PingFederate now ignores the encryption certificate if any is included in the metadata.

PingFederate 8.3.3 - May 2017

PingFederate® 8.3.3 is a cumulative maintenance release for PingFederate 8.3, which introduced many new features, such as self-service password reset (SSPR), hardware security module (HSM) hybrid mode, access token and authentication session sync, and improvements in OpenID Connect Relying Party support and OAuth Authorization Server. For more information, see the [release notes for PingFederate 8.3](#).

Resolved issues

Ticket ID	Description
PF-15695	The CIDR Authentication Selector no longer returns an error when processing a request in which the source IP address includes a port number.
PF-15682	Improved self-service password reset (SSPR) for the release of PingID® authentication policies.
PF-15628	Updated the configuration file for the PingFederate Windows service to improve the responsiveness of the logging service.
PF-15581	Resolved an issue where a memory leak could occur when IdP connections were configured with SLO profiles or authentication sessions were enabled.
PF-15571	The Account Locking Service (for account lockout protection) now tracks the number of sign-on failures when resource owner credentials are validated against multiple LDAP Username Password Credential Validator mappings on the OAuth Settings > Resource Owner Credentials Mappings screen.
PF-15472	Fixed a validation error where the <code>/authenticationPolicies/default</code> administrative API endpoint would reject a policy update that contained a reference to an IdP connection.
PF-15444	The <code>/pf-scim/v1/Users</code> and <code>/pf-scim/v1/Groups</code> SCIM inbound provisioning endpoints now return a 404 HTTP status code when the specified resource (a user or a group) is not found. (For more information, see the SCIM specification at www.simplecloud.info/specs/draft-scim-api-01.html#anchor6 .)
PF-15205	When using PingID for SSPR, PingFederate now uses the regional information found in the <code>pingid.properties</code> file.
PF-15204	Resolved an issue where PingFederate could not load connections with partner's certificates after it was integrated with a Luna SA HSM.
PF-14956	Resolved an issue where PingFederate could not process encrypted SAML messages when it was integrated with an nShield Connect HSM.

PingFederate 8.3.2 - March 2017

PingFederate® 8.3.2 is a cumulative maintenance release for PingFederate 8.3, which introduced many new features, such as self-service password reset (SSPR), hardware security module (HSM) hybrid mode, access token and authentication session sync, and improvements in OpenID Connect Relying Party support and OAuth Authorization Server. For more information, see the [release notes for PingFederate 8.3](#).

Resolved issues

Ticket ID	Description
PF-15233 and PF-15208	Addressed two potential security vulnerabilities when using PingFederate with PingDirectory™, formerly known as UnboundID Data Store (SECNT007).
PF-15120	PingFederate installer for Windows now detects the installation directory of Oracle Server JRE 8 update 121.
PF-14996	When the Change Password Email Notification option was enabled on an HTML Form Adapter instance, PingFederate sent the change password notifications to users at their email address stored in the mail attribute regardless whether another attribute was configured in the corresponding LDAP Username Password Credential Validator instance.
PF-14976	Fixed an issue where PingFederate did not write AUTHN_ATTEMPT events to the security audit log stored on an Oracle MySQL database.
PF-14972	Enhanced replay prevention on the resume URL when the authentication process involves user consent.
PF-14896	Resolved an issue where PingFederate did not update the Issuer field value in an IdP connection after loading the metadata from the OpenID Connect Provider.
PF-14877	Fixed an administrative console issue where administrators may see OpenID Connect as an available option on the Connection Type screen when creating an SP connection.
PF-14843	Self-service password reset (if configured) now works immediately after LDAPS is enabled on the underlying data store.
PF-14834	Fixed an issue where PingFederate could not process SAML authentication request messages that contain the Scoping element without the IDPList child element.

PingFederate 8.3.1 - January 2017

PingFederate® 8.3.1 is a cumulative maintenance release for PingFederate 8.3, which introduced many new features, such as self-service password reset (SSPR), hardware security module (HSM) hybrid mode, access token and authentication session sync, and improvements in OpenID Connect Relying Party support and OAuth Authorization Server. For more information, see the [release notes for PingFederate 8.3](#).

Resolved issues

Ticket ID	Description
PF-14819	New passwords submitted through self-service password reset are now only validated against the password policy of the associated data store, not the PingFederate internal service credentials password policy.
PF-14780	The outbound provisioning CLI (provmgr.bat) failed to start in a Microsoft Windows environment.

Ticket ID	Description
PF-14778	Resolved an issue where the administrative API could not create or update a connection with OGNL expressions.
PF-14771	PingFederate can now connect to the JSON Web Key Set endpoint (/pa/oidc/JWKS) on PingAccess.
PF-14631	Fixed an issue with the handling of client certificates that were set in HTTP headers by an HTTPS-terminating front end proxy server (SECNT005).
PF-14627	Upgraded to jose4j 0.5.3. PingFederate can now issue signed and encrypted JWT-based access tokens in an environment where PingFederate is integrated with a Thales HSM.
PF-14567	For any SLO-enabled SAML 2.0 IdP connection with multiple virtual server IDs, PingFederate now sends one logout request to the associated IdP when it receives a logout request at its SLO application endpoint (/sp/startSLO.ping).
PF-14423	Resolved an issue where the OAuth introspection endpoint (/as/introspect.oauth2) could not process validation requests when persistent grants were stored on an LDAP server.
PF-14372	PingFederate now includes SLO endpoint information in its SAML metadata.
PF-13996	Fixed an issue where PingFederate did not URL-encode the percentage character (if any was used in the Partner's Entity ID field of a connection) when constructing the SSO Application Endpoint URL on the connection summary screen.
PF-13624	When using OGNL expressions to define issuance criteria for authentication policy contract mapping, administrators may now reference attributes from an IdP connection using its Partner's Entity ID value; for example:

```

...
#idpCxnUsername = #this.get("idp.https://
sso.idp.local.SAML_SUBJECT")
...

```

In this example, the **Partner's Entity ID** value of the IdP connection is `https://sso.idp.local` and the referenced attribute is `SAML_SUBJECT`.

PingFederate 8.3 - December 2016

Enhancements

Self-service password reset

PingFederate® now offers self-service password reset (SSPR) for end users to recover their accounts in the event of forgotten passwords. Integrated into the HTML Form Adapter and Password Credential Validator (PCV) framework, end users can now reset their passwords via one of four different mechanisms:

- One-time link via email
- One-time password via email
- One-time password via text message
- PingID®

The LDAP Username PCV supports the SSPR features, relying on additional configuration to look up the required attributes for the chosen reset mechanism. Custom PCV implementations may also be developed to offer SSPR features for users stored in non-LDAP data sources.

Hardware security module hybrid mode

Configuring PingFederate to use a hardware security module (HSM) is now simpler with a new hybrid mode option. HSM hybrid mode allows cryptographic material to be stored either locally on the file system or on an attached HSM (such as SafeNet or Thales). At creation or import time, administrators can choose where to store the certificate and key material. This feature can be used as part of a transition strategy to an HSM-only configuration for optimal security in your deployment.

Access token and authentication session sync

To provide end users with a consistent experience across web and API based native applications and single-page applications (SPAs), the validation of OAuth access tokens and authentication sessions can now be linked together.

Enabled through a policy on Access Token Manager instances, the validation of the access token can depend on an associated authentication session from which it was authorized. Activity on the access token, triggered during a validation or introspection call to PingFederate, can subsequently update the activity on the authentication session thereby centrally managing idle and maximum lifetime session controls. Additionally, should an end user sign off and revoke the associated authentication session, access tokens that rely on it can be considered invalid.

This capability can be used in conjunction with PingAccess[®] to support its authorization controls on both web and API resources.

OpenID Connect Relying Party improvements

Relying Party support has been enhanced for improved interoperability and ease of administration.

Form POST and **Form POST with Access Token** login types have been added to complement the existing Authorization Code Flow support to provide a more optimized login process.

Custom request parameters can now be added to provide additional partner specific parameters during login. As an example, this may be used for integrating Google login services to provide its hd parameter (see <https://developers.google.com/identity/protocols/OpenIDConnect#hd-param>).

Third party initiated login is now supported as defined in the OpenID Connect standard (see http://openid.net/specs/openid-connect-core-1_0.html#ThirdPartyInitiatedLogin).

Additionally, the administrative API has been extended to include the management of OpenID Connect Browser SSO IdP connections.

OAuth Authorization Server enhancements

OAuth Authorization Server capabilities have been enhanced for greater deployment flexibility and reduced administrative overhead.

The storage of clients can now be customized using a new `ClientStorageManager` interface available in the PingFederate SDK. This provides another deployment option for managing clients in addition to the existing options of local XML configuration files and a relational database system.

Access token validation has been simplified to ease application integration. Resource server clients can be configured to allow validation to occur against all eligible Access Token Manager instances automatically. When enabled, resource server clients no longer need to specify additional parameters (such as `access_token_manager_id` and `aud`) to disambiguate the Access Token Manager instance. This feature can also simplify interactions with PingAccess by avoiding the need to align resource URIs between it and PingFederate.

Registered claims (such as `exp`) within JSON Web Tokens formatted access tokens can now be overridden within attribute fulfillment. If the `aud` claim must be fulfilled at runtime, instead of relying on the static **Audience Claim Value** configuration setting, the `aud` will be ignored during access token validation.

Other improvements

The following bundled components and third-party dependencies have been updated:

- Jetty 9.3.14 (SSLv2Hello disabled)
- jose4j 0.5.2

Resolved issues

Ticket ID	Description
PF-14397	Upgraded to Jetty 9.3.14 to resolve a rare error condition (<code>java.lang.StackOverflowError</code>).
PF-14369	OAuth clients that authenticate through IdP connections no longer fail to do so when message customizations have been configured in such connections.
PF-14359	PingFederate now sets the Accept HTTP request header to <code>application/json</code> in its SCIM outbound requests to delete groups.
PF-14222	When an SP adapter instance was mapped to a URL on the Target URL Mapping screen, PingFederate always failed to locate the SP adapter instance if the adapter instance was overridden within an IdP connection.
PF-14178	The SSO Directory Service excluded OpenID Connect (OIDC) IdP connections from its responses.
PF-14172	<p>The <code>/authenticationPolicies/default</code> administrative API endpoint failed to create or update a policy when the policy included at least one <code>attributeRules</code> with <code>"fallbackToSuccess": false</code>—the equivalent of creating or updating a policy in the administrative console with at least one rule, in which the Default to Success option was disabled; for example:</p> <pre data-bbox="537 911 1455 1226"> ... "attributeRules": { "items": [{ ... }], "fallbackToSuccess": false } ... </pre>
PF-14147	When an administrator deletes one or more partner protocol endpoints from a connection, the administrative console now removes the intended entries.
PF-14141	OIDC Relying Party support has been enhanced to handle multi-valued attributes.
PF-14113	Common Domain Service (for IdP Discovery) now works with the <code>pf.runtime.context.path</code> property in the <code>run.properties</code> file.
PF-14098	<p>The <code>/idp/adapters</code> and <code>/idp/adapters/{id}</code> administrative API endpoints failed to create or update IdP adapter instance when the adapter instance included at least one <code>attributeSources</code> lookup with <code>"memberOfNestedGroup": true</code>—the equivalent of creating or updating an IdP adapter instance in the administrative console with at least one LDAP data store lookup to query the memberOf attribute with the Nested Groups option; for example:</p> <pre data-bbox="537 1661 1450 1925"> ... "attributeMapping": { "attributeSources": [{ "type": "LDAP", ... "memberOfNestedGroup": true }], </pre>

Ticket ID	Description
	<pre> ... } ... </pre>
PF-14097	When loading metadata from an OIDC partner in an IdP connection, if the partner's HTTPS endpoint requires a certificate to be imported via the Trusted CAs screen, the certificate will now be trusted immediately after the import without requiring a restart.
PF-14083	SP-initiated SLO no longer fails if the <code>pf.runtime.context.path</code> property is set in the <code>run.properties</code> file.
PF-14029	The administrative console now prevents administrators from removing a JDBC data store that is used by one or more authentication policy contracts.
PF-13968	<p>The PingFedeate IdP server used to reject signed authentication requests when all the following conditions were met:</p> <ul style="list-style-type: none"> • The SP-initiated SSO and SP-initiated SLO profiles were enabled. • The Require digitally signed AuthN requests option was disabled. • The SP used different signing certificates to sign its authentication requests and SLO messages.
	PingFederate now logs a warning message to the server log instead.
PF-13967	In a clustered PingFederate environment, an engine node now merges the new keys obtained from the console node during a configuration replication process into its <code>pf.jwk</code> file.
PF-13952	The Windows service file (<code>PingFederateService.conf</code>) has been improved to resolve a compatibility issue between the Thales nShield client driver and Oracle Java 8.
PF-13824	When PingFederate was configured with a managed SP connection to PingOne® and such connection was subsequently used as the source to create a new connection (via the Copy feature on the SP Connections screen), the connection to PingOne would break. As a result, users could not sign on to PingOne dock.
PF-13622	Improved the Composite Adapter to handle null attribute values in its attribute mapping process.
PF-13522	The HTML Form Adapter's default sign-on template file (<code>html.form.login.template.html</code>) now renders properly when invoked by Microsoft Office 365 applications.
PF-12734	The <code>/dataStores</code> administrative API endpoint does not require an <code>id</code> value anymore to create a new data store. A system-generated, read-only value will be assigned instead.
PF-12489	If an IdP connection is configured with multiple virtual server IDs, the <code>AudienceRestriction</code> 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.

PingFederate 8.2.2 - October 2016

PingFederate® 8.2.2 is a cumulative maintenance release for PingFederate 8.2, which introduced many new features, such as OpenID Connect Relying Party support, authentication session caching, and password management improvements. For more information, see the [release notes for PingFederate 8.2](#).

Resolved issues

Ticket ID	Description
PF-14033	Increased the maximum allowed size for an HTTP request header to handle requests with large headers (for example, a request with a Kerberos ticket whose user belongs to many groups).
PF-14017	Addressed an issue where the LDAP Username Password Credential Validator requested an LDAP server to sort a result set with one result, which could inadvertently cause a delay in receiving a response from the LDAP server.
PF-13994	Improved the responsiveness of the administrative console when storing OAuth clients in database.
PF-13982	Resolved an issue where the Kerberos Adapter and the Kerberos Token Processor could not extract the SID attribute value from Kerberos tickets.

PingFederate 8.2.1 - September 2016

PingFederate 8.2.1 is a cumulative maintenance release for PingFederate 8.2, which introduced many new features, such as OpenID Connect Relying Party support, authentication session caching, and password management improvements. For more information, see the [release notes for PingFederate 8.2](#).

Resolved issues

Ticket ID	Description
PF-13950	PingFederate now redirects requests properly based on the HTTP Header for Hostname setting on the Server Configuration > Server Settings > System Options screen.
PF-13948	Resolved an issue where PingFederate finished its startup process without starting the integrated RADIUS server (part of the bundled PingID PCV).
PF-13928	PingFederate only returned a partial list of attributes for a selected object.
PF-13905	Resolved a logging issue that could lead to excessive memory usage.
PF-13837	Resolved an issue where attribute values were not consistently masked in the server log despite a proper configuration.
PF-13836	The upgrade process no longer copies older program files (from the previous releases of the PingID integration kit) to the new PingFederate 8.2.1 installation.
PF-13835	Improved the account linking process to handle the creation of new users without generating an <code>Unexpected exception</code> error condition. Note that new users were created despite the error message.
PF-13829	Improved PingFederate to validate OCSP responses signed using RSA SHA256, SHA384, or SHA512.
PF-13828	Enhanced the administrative console to handle the scenario where it cannot establish a JDBC connection with the database server that has been configured to store OAuth clients.
PF-13822	Administrative users without the User Admin role can now change their own passwords on the Server Configuration > Account Management screen.
PF-13740	Reverted the name of a parameter for the <code>/serverSettings/notifications</code> administrative API endpoint to <code>notifyAdminUserPasswordChanges</code> . Administrative

Ticket ID	Description
PF-13730	API calls written for PingFederate 8.1.4 (and earlier versions) are now compatible with version 8.2.1.
PF-13729	Resolved an issue where PingFederate did not include the event type (OAuth), the subject, and the protocol (OIDC) in the audit information for requests sent to the UserInfo endpoint.
PF-13706	When IdP connections were used as part of the authentication policies, PingFederate failed to write information to the audit log after the users had finished with such IdP connections.
PF-13670	Updated the bundled PingID Integration Kit to version 1.3.3.
PF-13670	When validating redirection URLs, the matching algorithm has been improved when evaluating the entries defined on the Server Configuration > Redirect Validation screen. A more specific match is considered a better match and an exact match is considered the best match.

PingFederate 8.2 - August 2016

Enhancements

OpenID Connect Relying Party support

The OpenID Connect protocol is now available for browser-based SSO when PingFederate® is acting as a service provider - a *Relying Party* in OpenID Connect terminology. This support complements the pre-existing OpenID Provider (OP) functionality and enables customers to make greater use of this modern protocol for partner use cases.

OpenID Connect can be configured using IdP connections and mapped into last mile integration options (such as OpenToken or Agentless Integration Kits) to support SSO into existing applications. These connections can also be bridged to other SSO protocols to act as a federation hub.

Authentication session caching

PingFederate now offers the ability to cache authentication sessions, giving administrators control over how often end users need to reauthenticate with any IdP adapter or IdP partner. Authentication session policies can be defined both globally and for individual authentication sources.

This feature can also be used to create a centralized session management experience for end users accessing protected applications through the Ping Identity Federated Access Management solution. PingAccess® web sessions will synchronize the session's active state at PingFederate via calls to the OpenID Connect Session Revocation API.

Authentication policy contract mapping for OAuth

Authentication policy contract support has been added to OAuth and OpenID Provider configuration, enabling administrators to centrally define authentication logic and attribute mapping that can be reused across many varying partners and protocols.

As part of the Federated Access Management solution, these capabilities enhance the contextual access management capabilities when deployed alongside PingAccess and PingID®.

Password management improvements

New password management capabilities have been added to improve usability and end-user security.

Optionally enabled in the HTML Form Adapter, end users can be warned during login about an upcoming password expiry. This allows end users to better manage when and where they perform a password update to balance security policy compliance with user experience.

Additionally, when end users update their password through the HTML Form Adapter, they can be emailed a notification about this event. End users suspicious about this event can then take action to ensure their account remains secure.

Improved SLO support

To improve security and the end-user's experience, Single Logout (SLO) support has been improved to cover more use cases and protocol mappings. When SLO is initiated, PingFederate sends SLO requests to the applicable partners associated with the user's session and clean up any local sessions.

The following configuration models now support SLO:

- Federation hub
- Adapter-to-adapter mappings
- OAuth and OpenID Connect mappings using Browser SSO IdP connections

Authentication context overrides

When interacting with an IdP partner, PingFederate now offers improved control over how authentication contexts are communicated with it. Available as an override on IdP connections, local authentication context values can be mapped into remote values to dictate how they are translated. This provides greater flexibility when working with IdP partners that have fixed names for authentication mechanisms that may differ from your desired local representation.

OAuth enhancements

OAuth Authorization Server capabilities have been enhanced with the following features:

- Support for OAuth 2.0 Token Introspection (as defined in [RFC7662](#)) has been added to provide a standards-based mechanism to validate tokens.
- Encryption can now be configured on JSON Web Token (JWT) formatted OAuth access tokens.
- The error responses for the Resource Owner Credentials grant type now provide more meaningful `error_description` parameter details. Messages can additionally be localized and customized through the PingFederate localization framework.
- OAuth Persistent Grant Expiry can now be defined in minutes.
- Elliptic Curve Digital Signature Algorithm (ECDSA) is now available for JWT formatted access tokens.
- The JSON Web Token Access Manager can now be configured to publish a JSON Web Key Set with the keys/certificates that can be used for signature verification of access tokens.

Administrative API enhancements

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

- Authentication selectors
- Authentication policies
- OAuth attribute mapping in IdP connections
- SAML 2.0 Bearer OAuth grant type mapping

Other improvements

- Inbound provisioning has been enhanced to support sorting and pagination of resources as defined in the SCIM 1.1 standard.
- Audit logging enhancements to provide better visibility into failed authentication attempts.
- The HTTP Request Parameter Selector now supports wildcard based matching.
- The license update process is now simpler and can be applied to a cluster during configuration replication.
- Performance improvements to both OpenID Connect Policy administration and OAuth flows.
- The Log4j 2 configuration has been optimized for performance by defaulting to INFO level and disabling console logging.

- The Active Directory and Kerberos configuration now supports the definition of fully qualified domain names for key distribution centers.
- The domain of the PF cookie is now configurable to facilitate sharing of the user's session amongst engine nodes using differing hostnames under a common domain.
- Email notification templates now support HTML formatting.
- The default user-facing screens have been redesigned.
- Various administrative console improvements.
- The PingID Integration Kit 1.3.1 is now bundled with PingFederate.
- The following bundled components and third-party dependencies have been updated:
 - OpenToken Adapter 2.5.5
 - Log4j 2 2.6.1

Resolved issues

Ticket ID	Description
PF-13588	The <code>soLingerTime</code> property of the secondary HTTPS connector (<code>httpsSecondaryConnector</code>) now uses the correct property name (<code>http.soLingerTime</code>) in the <code><pf_install>/pingfederate/etc/jetty-runtime.xml</code> file.
PF-13585	PingFederate now verifies database connections if a SQL statement (the Validate Connection SQL field value) is provided in the JDBC data store configuration.
PF-13336	The Attribute Sources & User Lookup configuration does not report an error anymore when the <code>tokenGroups AD user attribute</code> is configured as a binary attribute in the LDAP data store configuration.
PF-13335	The OAuth token endpoint (<code>/as/token.oauth2</code>) has been enhanced to return the HTTP status code of 500 when a token request is interrupted by an unexpected error, such as a data source failure or a runtime exception.
PF-13202	Resolved an issue where a request did not failover gracefully from the Kerberos Adapter through an authentication policy.
PF-13048	The attempts to change password (through the HTML Form Adapter) from locked out users no longer trigger PingFederate to write the Unexpected Runtime Authn Adapter Integration Problem error message in the server log. Note that the user experience and the audit log behavior remain the same: <ul style="list-style-type: none"> • The attempts are rejected with this message: Your account is locked. Please contact your system administrator. • The attempts are recorded to the audit log.
PF-12906	Improved a shell script (<code><pf_install>/pingfederate/sbin/pingfederate-run.sh</code>) that the PingFederate install script uses.
PF-12751	When creating or updating a connection, the administrative API no longer requires the <code>validFrom</code> and <code>expires</code> properties of a certificate to be accurate down to the millisecond. Both of the following snippets are now accepted: <ul style="list-style-type: none"> • Without milliseconds:

```
...
"validFrom": "2013-05-16T10:35:39Z",
"expires": "2056-05-10T10:35:39Z",
...
```


Ticket ID	Description
	<ul style="list-style-type: none"> With milliseconds: <pre data-bbox="578 247 1463 390">... "validFrom": "2013-05-16T10:35:39.000Z", "expires": "2056-05-10T10:35:39.000Z", ...</pre>
PF-12709	PingFederate no longer writes an <code>ERROR</code> message with a stack trace when a request results in a failure through an authentication policy. The simplified information is categorized as <code>DEBUG</code> information now.
PF-12648	<p>The Server Configuration > Account Management screen is only accessible when one of the following conditions is met:</p> <ul style="list-style-type: none"> The administrative console is configured to use native authentication. The administrative API is configured to use native authentication. Both the administrative console and the administrative API are configured to use native authentication (the default as specified in the <code><pf_install>/pingfederate/bin/run.properties</code> file).
PF-12378	The Server Configuration > Certificate Management configuration no longer results in an error when PingFederate encounters a connectivity issue with the OAuth client data store.
PF-12364	PingFederate now properly handles the error condition where the <code>wa=wsignoutcleanup1.0</code> query parameter is incorrectly sent to the <code>/idp/prp.wsf</code> endpoint.
PF-12325	Resolved an issue where OAuth parameters <code>client_id</code> and <code>client_secret</code> were accepted as URL query parameters.
PF-12274	Improved the <code><pf_install>/pingfederate/server/default/conf/terse.example.log4j2.xml</code> file; for example, audit information is now preserved.
PF-12136	Resolved an issue where the Kerberos Adapter and the Kerberos Token Processor could not extract the SID attribute value when the Kerberos environment was configured to use AES encryption.
PF-12131	PingFederate does not write an <code>ERROR</code> message with a stack trace anymore when the Kerberos Adapter receives an NTLM token (as opposed to a valid Kerberos ticket). The simplified information is categorized as <code>DEBUG</code> information now.
PF-11920	When issuing JWT formatted OAuth access tokens based on refresh tokens, PingFederate did not include the <code>client_id</code> information in the access tokens.
PF-11411	Resolved an issue where the outbound provisioning threads were closing after 60 seconds.
PF-11209	When a WS-Trust STS transaction completes successfully using the Kerberos Token Processor, PingFederate now writes the user identifier to the audit log.

PingFederate 8.1.4 - July 2016

PingFederate 8.1.4 is a cumulative maintenance release for PingFederate 8.1, which introduced many new features, such as advanced authentication policies, elastic scaling compatibility, and metadata publishing and consuming. For more information, see the [release notes for PingFederate 8.1](#).

Resolved issues

Ticket ID	Description
PF-13353	The Crypto administrative role is no longer required to create a connection; only the Admin administrative role is required.
PF-13271	Deploying a configuration archive file to the <pf_install>/pingfederate/server/default/data/drop-in-deployer directory no longer triggers a certificate error that may block access to the PingFederate administrative console via LDAP authentication.
PF-13215	Resolved a compatibility issue in the LDIF scripts designed for UnboundID Data Store. The new scripts can be found in the <pf_install>/pingfederate/server/default/conf/access-grant/ldif-scripts directory.
PF-12787	Resolved an issue where PingFederate might stop responding after a non-blocking socket operation could not be completed.
PF-12720	In rare occasions, PingFederate installer for Windows might fail to detect the progress of the upgrade and report the upgrade as succeeded.
PF-12683	The PingFederate install script now offers the opportunity to update PingFederate's JAVA_HOME environment when the user's JAVA_HOME environment points to a different location.
PF-12682	PingFederate omitted the CN of expiring certificates in server log messages.
PF-12491	PingFederate ignored the default target URL of an adapter-to-adapter mapping when a request was initiated with both the IdpAdapterId and SpSessionAuthnAdapterId parameters and no default SSO URL was specified on the SP Configuration > Default URLs screen.
PF-12481	An end user might encounter a redirect loop after initiating an SLO request.
PF-12469	Improved the hsmypass and obfuscate command-line utilities to handle hotfix packages.
PF-11683	Improved the redirect validation process for the InErrorResource parameter.

PingFederate 8.1.3 - May 2016

PingFederate 8.1.3 is a cumulative maintenance release for PingFederate 8.1, which introduced many new features, such as advanced authentication policies, elastic scaling compatibility, and metadata publishing and consuming. For more information, see the [release notes for PingFederate 8.1](#).

Resolved issues

Ticket ID	Description
PF-12473	Editing the configuration of a token translator instance that is currently in-use by one or more WS-Trust requests no longer results in a deadlock.
PF-12468	Improved the handling of simultaneous requests from the same origin.
PF-12457 and PF-12454	Resolved an access issue in the /ConfigArchive administrative API endpoint and the Server Configuration > Configuration Archive screen in the administrative console.
PF-12450	The administrator audit log (admin.log) has been enhanced to include activities from configuration archive import.
PF-12412	Improved the Kerberos Adapter to resolve a browser looping issue when the following conditions were met:

Ticket ID	Description
	<ul style="list-style-type: none"> • PingFederate engine nodes (in a clustered PingFederate environment) were deployed behind a load balancer or a reverse proxy server without sticky sessions. • The authentication flow involved an instance of the Kerberos Adapter that was part of an instance of the Composite Adapter.
PF-12392	OAuth persistent grants that are stored in an LDAP directory server are now removed as the grants expire.
PF-12375	Custom data store filters could not be modified.
PF-12300	In a clustered PingFederate environment, the <code>/pf/JWKS</code> endpoint from the engine nodes might serve different keys for the same key ID (<code>kid</code>).
PF-12135 and PF-12134	Enhanced PingFederate to support surrogate pair characters in the administrative console and for runtime events.

PingFederate 8.1.2 - April 2016

PingFederate 8.1.2 is a cumulative maintenance release for PingFederate 8.1, which introduced many new features, such as advanced authentication policies, elastic scaling compatibility, and metadata publishing and consuming. For more information, see the [release notes for PingFederate 8.1](#).

Resolved issues

Ticket ID	Description
PF-12383	Resolved a logging issue that could lead to degraded performance when an error had occurred.
PF-12323	Extended contract is now available for all IdP adapters.
PF-12322	When the Adapter-to-Adapter Mappings configuration contains one or more mappings from only one IdP adapter instance, PingFederate no longer returns a <i>Sign On Error</i> message when an adapter-to-adapter request provides an <code>IdpAdapterId</code> query parameter without a value.
PF-12289	Resolved a security issue in the administrative console.
PF-12285	PingFederate did not issue OAuth refresh tokens when the Roll Refresh Token Values (default policy) check box was selected and the Minimum Interval to Roll Refresh Tokens (hours) field was set to 0 in the OAuth Settings > Authorization Server Settings screen.
PF-12283	The LDAP query process has been improved to handle result sets that may contain a large number of objects.
PF-12238	Resolved a logging issue when PingFederate was deployed with Box Connector 1.0.
PF-12203	The Asynchronous Front-Channel Logout process has been enhanced to terminate sessions at the grant-management endpoint (<code>/as/oauth_access_grants.ping</code>).
PF-12182	PingFederate now supports Proof Key for Code Exchange by OAuth Public Clients (tools.ietf.org/html/rfc7636) through a new parameter (<code>code_challenge_method</code>) at the OAuth authorization endpoint (<code>/as/authorization.oauth2</code>).
PF-12158	Resolved an issue where LDAP schema caching might lead to a high usage of memory.

PingFederate 8.1.1 - February 2016

PingFederate 8.1.1 is a cumulative maintenance release for PingFederate 8.1, which introduced many new features, such as advanced authentication policies, elastic scaling compatibility, and metadata publishing and consuming. For more information, see the [release notes for PingFederate 8.1](#).

Resolved issues

Ticket ID	Description
PF-12237	The Kerberos Token Adapter and Kerberos Token Processor now return a null value for the <code>SID</code> attribute when such attribute cannot be decrypted.
PF-12235	The policy engine now skips authentication policies that only contain closed-ended failed paths when processing a request.
PF-12190	The PingFederate administrative console no longer allows IdP connections that have been placed in one or more authentication policies to be deleted from the IdP Connections screen. For each of these in-use IdP connections, administrators can click Check Usage to determine which policy (or policies) must be reconfigured before the IdP connection can be removed.
PF-12142	In the OAuth Settings > Client Management screen, pressing the Enter key in the search field no longer returns the administrator to the OAuth Settings menu.
PF-12141	The OAuth Settings > Authorization Server Settings screen now displays scope and scope group values in alphabetical order.
PF-12133	The optional Enable 'Remember My Username' feature in the HTML Form Adapter now works with the latest Google Chrome (version 48).
PF-12080	The maximum lifetime value of an access token in JSON Web Token format was incorrect.
PF-11601	Updated the time stamp format in <code>log4j2.xml</code> file for better compatibility with Oracle Database.
PF-11444	PingFederate now provides more information in the server log when the SCIM inbound provisioner fails to provision users to Active Directory.

PingFederate 8.1 - January 2016

Enhancements

Advanced authentication policies

PingFederate now offers the ability to define central authentication policies that chain together selectors, adapters, and IdP connections to control how end users login for SSO and OAuth use cases. Additionally, data source lookups can now be performed within IdP adapter instances to associate them more closely to the authentication source and make policy decisions based on their results.

As part of the Federated Access Management solution, these capabilities support context-sensitive adaptive authentication when combined with PingAccess and PingID®. Identity data following a first-factor of authentication can be augmented to make intelligent decisions on how to further authenticate a user. A common scenario for this could be group-based selection of multifactor authentication.

Metadata publishing and consuming

Browser SSO connection creation and ongoing administration is now simpler with SAML metadata publishing and consuming. This feature provides support for metadata to be retrieved from and published to HTTP or HTTPS URLs and compatibility with trust networks such as InCommon.

Partner connections can be associated with metadata URLs, and a policy provides controls over periodic checks of the metadata to automatically update certificates and contact information. If other connection changes are detected within the metadata, the PingFederate administrators can be notified via email of said changes to review for manual reconciliation. Additionally, certificate and key rotation features enable automated rollover of SAML signing and encryption certificates for partners that can consume metadata URLs.

Elastic scaling compatibility

PingFederate now provides improved compatibility with auto (or "elastic") scaling features provided by cloud infrastructure services including Amazon Web Services and OpenStack. The dynamic discovery mechanism enables new members to join an existing cluster without configuring IP addresses ahead of time; it enables the horizontal scaling of a cluster without manual intervention when traffic volume dictates additional resources are required.

LDAP directory and custom storage for OAuth persistent grants

For OAuth deployments requiring persistent grants, this data can now be stored in an LDAP directory as an alternative to a relational database system. Consult [System requirements](#) in the *Get started with PingFederate* guide for the list of qualified directory server products for this feature.

Additionally, PingFederate now allows developers to build their own OAuth Access Grant Manager, providing more options (beyond an LDAP directory server or relational database system) for how their access grant can be persisted. The PingFederate SDK provides the `com.pingidentity.sdk.accessgrant.AccessGrantManager` interface that customers can implement. A sample implementation is located in the `<pf_install>/pingfederate/sdk/plugin-src/access-grant-example/java` directory.

Active Directory Authentication Mechanism Assurance support

The Kerberos Adapter has been extended to return the SID (security identifier) from the Kerberos ticket as part of its contract. This new capability enables the checking of the user's Authentication Mechanism Assurance to determine if the user logged in to the Windows desktop via username/password or a smart card. Authorization policies, either within PingFederate or the target application, can then dictate what the user is allowed to do.

PingFederate installer improvements

Windows installer and Red Hat Enterprise Linux install script can now detect that a previous release is present on the system that was also installed using an earlier release of the installer and automate the upgrade.

Administrative API enhancements

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

- SAML 2.0 Browser SSO connections using Artifact or SOAP bindings and SLO
- SAML 1.x connections
- WS-Federation connections
- Target URL mappings
- Metadata URLs management
- Key pair and certificate rotation for self-signed certificates
- Authentication policy contracts
- Authentication policy contract mapping in IdP and SP connections

Security enhancement

- XML encryption features within SAML Browser SSO connections have been extended to support a primary XML decryption key and secondary XML decryption key.

Other improvements

- The default user-facing screens have been beautifully redesigned. Previously modified screens can be preserved by copying the respective Velocity template files from the old installation to the new installation.

For more information, see [Copy customized files or settings](#) on page 630 in the *PingFederate Upgrade Utility user guide*.

- Kerberos Token Processor is now bundled with PingFederate.
- CIDR Authentication Selector now supports IPv6 addresses.
- UnboundID Data Store is now qualified for various directory requiring features.
- Oracle database 12c is now qualified.
- OAuth client ID is now available as a variable in applicable end user facing Velocity templates.
- Connection Mapping Contract has been renamed to Authentication Policy Contract.
- Better administrator experience when data stores are unresponsive.
- Capability to map the SAML 2.0 `Authenticating Authority` attribute value to an SP adapter contract.
- Upgraded to JGroups 3.6.5 (and JGroup MERGE3).

Resolved issues

Ticket ID	Description
PF-12240	PingFederate did not capture the response size in the request log.
PF-11817	The PingFederate administrative APIs did not save the contract fulfillment mappings when creating an SP connection with a pseudonym identity mapping and additional attributes from multiple data stores using one mapping.
PF-11685	PingFederate logged the runtime port as -1 in the transaction log.
PF-11684	When the Enable 'Remember My Username' and Allow Password Changes check boxes were selected, and a user had successfully authenticated and saved the username, another user could not sign in using the same HTML Form Adapter instance from the same browser if a password change was required as well.
PF-11672	PingFederate could not provision users with default groups to the PingOne directory.
PF-11541	The PingFederate administrative console reported an error if it was configured to use LDAP authentication without a prior login using the default native authentication.
PF-11528	IdPs, who were invited by the customers of PingOne SSO for SaaS Apps, could not establish a managed SP connection to PingOne using the Connect to PingOne wizard.
PF-11510	Just-in-time (JIT) provisioning was not updating group membership information consistently.
PF-11469	Resolved SCIM inbound provisioning issues on Active Directory with partitions.
PF-11432	PingFederate logged non-fatal <code>Unexpected</code> exception messages to the transaction log when handling OpenID Connect requests.
PF-11359	The <code>configcopy</code> utility failed and returned a <code>No such operation 'saveServerFile'</code> message.
PF-11318	Added randomization support in the resume path for the Kerberos Adapter.
PF-11317	Resolved configuration issues in JDBC connection pooling.
PF-11175	Resolved an issue where PingFederate was unable to replicate provisioner configuration file when provisioner was running at the same time
PF-11130	Improved PingFederate Upgrade Utility to remind the administrators to migrate from an older version of Username Token Processor (if found) to the newer Username Token Processor bundled with PingFederate since version 7.2.

Ticket ID	Description
PF-11085	When the Allow Password Changes check box is selected, the HTML Form Adapter reauthenticates the user using the new password immediately after a successful password change request. PingFederate 8.1 introduced a new HTML Form Adapter setting, Post-Password Change Re-Authentication Delay , to handle the scenario where a time delay between the password change and the reauthentication attempt is required.
PF-11079	Previously configured OAuth access token mapping returned an error when the partner's entity ID of the corresponding IdP connection had changed.
PF-10882	Starting with PingFederate 8.1, you are not required to place an instance of the Requested AuthN Context Authentication Selector as the last checkpoint in a policy in order for the AuthN Context Attribute value to be mapped into an assertion. Similarly, you are also not required to place an instance of the CIDR Authentication Selector as the last checkpoint in a policy in order to use the Result Attribute Name value to fulfill an attribute contract or for issuance criteria.
PF-10811	When querying multiple data stores for contract fulfillment or token authorization, if a data store uses results from previous queries as input, and if the previous queries return no result, PingFederate should continue the SSO process by moving on to the next data store in the setup. For more information, see Attribute mapping with multiple data sources on page 506.
PF-10569	The metadata import process did not compare the bindings included in the metadata against those defined in the SAML specifications using a case-sensitive matching rule. As a result, requests might fail at runtime. This issue is now resolved; for example: <ul style="list-style-type: none"> • <code>urn:oasis:names:tc:SAML:2.0:bindings:HTTP-POST</code> is valid. • <code>urn:oasis:names:tc:SAML:2.0:bindings:HTTP-Post</code> is invalid because the binding ends with <code>HTTP-Post</code> as opposed to <code>HTTP-POST</code>.
PF-10149	For email notification using SSL or TLS, hostname verification of the certificate is now available. 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 PingFederate Upgrade Utility preserves the configuration without activating the hostname verification option for compatibility reasons. Administrators should consider activating this new option for greater security.

PingFederate 8.0.4 - November 2015

PingFederate 8.0.4 resolved the following issues.

Ticket ID	Description
PF-11463	Resolved a potential security vulnerability.
PF-11413	The <code>/pf-admin-api/v1/idp/spConnections</code> administrative API no longer throws a <code>NullPointerException</code> error after upgrade.
PF-11328	Upgraded the <code>jaxb-impl.jar</code> file from version 2.1.7 to version 2.2.14 to resolve a duplicate Java class error when the WAM Integration Kit was deployed with the library files from Oracle (for the use case of WAM plug-in for Oracle Access Manager).
PF-11287	When importing a metadata file, PingFederate no longer throws a <code>NullPointerException</code> error when the metadata file is invalid.

Ticket ID	Description
PF-11198	Resolved a duplicate Java class error when the WAM Integration Kit was deployed with the <code>oraclepki.jar</code> file from Oracle (for the use case of WAM plug-in for Oracle Access Manager).
PF-11161	The OAuth authorization endpoint (<code>/as/authorization.oauth2</code>) now returns a 403 Forbidden response when it receives an HTTP HEAD request.
PF-11146	The LDAP connection pool does not leak connections anymore after its associated data store settings have been modified.
PF-11106	PingFederate returned a 400 Bad Request response when SCIM inbound provisioning was configured with one or more custom SCIM attributes and the Nested Groups check box was selected for the <code>memberOf</code> attribute in the configuration for SCIM responses.
PF-11082	The <code>/pf-scim/v1/Groups</code> endpoint for SCIM inbound provisioning did not return the schema attribute.
PF-10911	PingFederate did not remove an attribute from the access token mapping configuration after such attribute had been removed from the access token contract.
PF-8624	End-user credentials for all WS-Federation transactions should always be masked in the server log.

PingFederate 8.0.3 - September 2015

PingFederate 8.0.3 resolved the following issues.

Ticket ID	Description
PF-10975	SSO and SLO requests using the SOAP binding failed with Oracle Java 8 Update 51.
PF-10963	Installation of the PingFederate Windows service (via the <code>install-service.bat</code> file) failed on Oracle Java 8.
PF-10894	A race condition could occur in the authentication process through the integrated Kerberos Adapter during high request volume.
PF-10878	An obfuscated password for cluster authentication on an engine node inadvertently caused PingFederate failed to start after an upgrade.
PF-10858	Despite the fact that an SSO request completed successfully, PingFederate logged a failure in the audit log when the optional Default Target URL was omitted in an IdP connection that was deployed as part of the federation hub use case.
PF-10846	Improved the HTTP responses returned by the UserInfo endpoint as follows: <ul style="list-style-type: none"> When the OpenID Connect role is not enabled, the UserInfo endpoint returns <code>404 Not Found</code>. When the OpenID Connect role is enabled but the request has failed the token authorization workflow (issuance criteria), the UserInfo endpoint returns <code>403 Forbidden</code> with an <code>WWW-Authenticate</code> header; the header value reads: <pre>error_description="ACCESS_DENIED" realm="Userinfo" scope="<a list of applicable scopes>"</pre> When the OpenID Connect role is enabled but the request has failed the attribute mapping process, the UserInfo endpoint also returns <code>403 Forbidden</code> with an <code>WWW-Authenticate</code> header; the header value reads: <pre>realm="Userinfo" scope="<a list of applicable scopes>"</pre>
PF-10810	Resolved a security issue in the administrative console. For more information, visit the Ping Identity Support website and refer to the Security Advisory SECADV011.

Ticket ID	Description
PF-10797	Improved the sample code for password credential validator (<code>SamplePasswordCredentialValidator.java</code>) to illustrate the use case better.
PF-10757	Enhanced <code>obfuscate.bat</code> and <code>obfuscate.sh</code> to support special characters. (Administrators must enclose passwords with special characters by single quotation marks.)
PF-10711	Provided support for boolean and integer values in JWT access token manager.
PF-10430	The inbound provisioning process used an incorrect naming context to search users, which might return an error in some Active Directory environments.
PF-10329	The <code>exclude-patterns</code> sample under <code>X-Frame-Options</code> in the <code><pf_install>/pingfederate/server/default/data/config-store/response-header-runtime-config.xml</code> file contained an incorrect sample.
PF-9839	When a signed <code>AuthnRequest</code> came with an <code>AssertionConsumerServiceURL</code> attribute, and the authentication failed, PingFederate sent the <code>SAMLResponse</code> (with a nonsuccess status) to the Assertion Consumer Service URL defined in the SP connection in error. PingFederate should send the nonsuccess status to the <code>AssertionConsumerServiceURL</code> value specified in the signed <code>AuthnRequest</code> .

PingFederate 8.0.2 - August 2015

PingFederate 8.0.2 resolved the following issues.

Ticket ID	Issue
PF-10657	The PingOne directory password credential validator fails to authenticate when PingFederate is deployed behind a web proxy server.
PF-10646	The Connect to PingOne wizard fails when PingFederate is deployed behind a web proxy server.
PF-10622	The PingFederate SP server does not provision attributes from SCIM Enterprise Schema Extension through the support of custom attributes in the SCIM inbound workflow.
PF-10616	Kerberos authentication fails unexpectedly after the initial login has succeeded.
PF-10553	Upgraded JGroups to version 3.6.4 to resolve a JGroups issue: <i>FD_SOCK/FD: members are not unsuspected</i> (JGRP-1922). For more information, see https://issues.jboss.org/browse/JGRP-1922 .

PingFederate 8.0.1 - July 2015

PingFederate 8.0.1 resolved the following issues.

Ticket ID	Issue
PF-10551	The PingOne directory password credential validator in PingFederate 8.0 does not support the change password feature for expired accounts. It should.
PF-10437	An upgraded PingFederate 8.0 may fail to start if a Jetty patch was applied to the older PingFederate installation.
PF-10388	The LDAP Username Password Credential Validator in PingFederate 8.0 does not include the DN attribute in its responses after the LDAP credentials have been validated. It should.

PingFederate 8.0 - June 2015

- PingOne integration improvements
- PingFederate installers for Microsoft Windows and Red Hat Enterprise Linux

- Administrative API enhancements; introduced the capability to authenticate by X.509 certificates, LDAP or RADIUS accounts and extended coverage to:
 - IdP and SP Default URLs configuration
 - Resource Owner Credentials Mappings (OAuth)
 - Kerberos Realms and Active Directory Domains configuration
 - Export/Import SP Connection metadata
- SCIM Provisioning enhancements to support:
 - Custom attributes for inbound and outbound provisioning
 - Filtering and querying on attributes for Inbound Provisioning
 - Nested group membership for Outbound Provisioning
- Update SAML connections using metadata XML files
- Browser SSO protocol customization
- JSON Web Token Support in WS-Federation SP connections
- HTTP Request Parameter Authentication Selector
- Kerberos IdP Adapter
- Logging enhancements:
 - Upgraded to Log4j 2
 - Introduced a new access token tracking ID
 - Introduced a new audit log for the administrative console with detailed information for each event
- Security enhancements
- Other enhancements:
 - Redesigned the administrative console experience
 - Enhanced compliance with OpenID Connect specification
 - Upgraded to Jetty 9.2.11
 - Removed the Oracle Java SE Development Kit (JDK) dependency in favor of the Oracle Java SE Runtime Environment (Server JRE)
 - Restructured context-sensitive help in the administrative console to hyperlink to online documentation

Resolved Issues

PingFederate 8.0 also resolved the following issues.

Ticket ID	Issue
PF-9916	Cluster nodes throwing RuntimeExceptions during RPC calls can fail operations for the whole cluster
PF-9905	The “Message” column in the server-log-sqlserver.sql script is not large enough
PF-9904	The “Description” column in the audit-log-sqlserver.sql script is not large enough
PF-9869	A challengeable PCV is not being re-challenged when it throws a PasswordCredentialChallengeException
PF-9838	An error occurs when setting the TargetResource to a URL that has multiple '#' chars
PF-9837	In an Office 365 deployment, Lync locks out an account if their password was changed in AD
PF-9733	LDAP's memberOf Nested Groups check box value is not saved when using memberOf in an OGNL expression
PF-9729	The Log4j 2 appender PatternLayoutForDB throws NPE for null messages (resolved by upgrading to Log4j 2)
PF-9565	Windows service installation script causes cosmetic error during auto-configuration of PingFederateService.conf

Ticket ID	Issue
PF-9549	Unable to create an SpConnection with multiple adapter mappings via the administrative API
PF-9476	The location value for IdpAdapterRef is incorrect in the administrative API's SpConnection
PF-9441	SOAP Binding does not output the fault details from the partner
PF-9359	Refreshing an OAuth token when a data store is not available revokes the access grant
PF-9355	MySQL access grant script fails if the sql mode is set to NO_ZERO_DATE
PF-9200	Japanese characters are not displayed correctly in http.error.page.template.html template
PF-9162	OAuth grant table gets queried even when refresh tokens/re-use of persistent grant are not being used
PF-9098	Username is not audited when using the pseudonym identity mapping
PF-9082	A connection's signing algorithm cannot be changed without a certificate change
PF-9057	LDAP attributes are not always masked in the log
PF-9043	Using Nested Groups does not escape special characters in the LDAP search filter
PF-9042	In an Office 365 deployment, Outlook locks out account if their password has been changed in AD
PF-9040	Connection Deployer does not validate binary attributes importing a connection
PF-9039	LDAP Username Password Credential Validator should not return all attributes for a DN query
PF-9004	“javax.servlet.http.NoBodyResponse cannot be cast to org.sourceid.servlet.HttpServletRespProxy” error when receiving HEAD requests
PF-8993	Overriding default trust manager should be logged at DEBUG
PF-8989	Creating OpenToken Adapter via the administrative API causes error at runtime
PF-8962	Deadlock can occur when using database logging with failover
PF-8683	RSTR is missing the RequestedAttachedReference for encrypted SAML tokens
PF-8575	Audit log shows 'success' when RST validation fails

PingFederate 7.3 - January 2015

- Federation Hub
- SampleAuthenticationSelector SDK Sample
- SP Connections API
- Group Support for Inbound Provisioning
- Support for wreply in WS-Federation SSO
- Improved Redirect Validation
- Improved the cleanup process of expired persistent grants; added an index (EXPIRESIDX) for the expires column in the pingfederate_access_grant database table
- Hostname verification for LDAPS data stores
- Issuer restriction in anchored trust model for back-channel authentication and XML signature verification
- Elliptic Curve algorithms for certificate creation and XML signatures
- RSA certificate creation enhanced to include RSA SHA-2 signing algorithms and 4096 bit keys
- Application Name and Application Icon URL for SP Connections and Adapter-to-Adapter mappings
- Customizable response from /pf/heartbeat.ping
- Customizable template for HTTP error pages
- Customizable favicon.ico

- LDAP Username Password Credential Validator can now be configured to return additional user attributes beyond the username and user DN
- PingOne Directory Password Credential Validator is now bundled as part of the standard PingFederate installation
- Key algorithm and key size are now displayed in certificate management pages and the certificate details popup
- HTML Form Adapter now supports a configurable maximum session lifetime
- Support for nested LDAP groups in Outbound Provisioning
- Support retrieval of nested LDAP groups in Attribute Contract Fulfillment
- Upgraded JGroups to version 3.5.1
- Upgraded Jetty to version 8.1.16
- Over 60 other product issues resolved

PingFederate 7.2 R2 - September 2014

- Multiple Access Token Management Plug-in Instances
- Improved Attribute Mapping for OAuth access-token contract
- Multiple data-source lookups for OAuth workflows
- OpenID Connect / OAuth 2.0 form POST response mode
- Support User Info in the OpenID Connect ID Token
- Administrative API enhancements:
 - Configuration archive management
 - IdP Connection Metadata Export/Import
 - Certificate Management of Trusted CAs, SSL Client Keys, and SSL Server Certificates
 - Data Sources
 - Password Credential Validators
 - Adapter-to-Adapter Mapping
- Administrative Console Enhancements
- Other Enhancements:
 - The OAuth client_id parameter is now available in the IdP Adapter SDK interface
 - Support SID encoding for binary LDAP attributes (enabling claims-based authentication to Microsoft Outlook Web Access)
 - Various security enhancements
 - Over 40 other product issues resolved

PingFederate 7.2.1 - August 2014

- Fixed an issue for Office 365 Outlook use cases where user accounts could be locked after changing their passwords in Active Directory.
- Fixed an issue where users would not be provisioned correctly under certain conditions.
- Fixed an issue that prevented additional "Actions" in adapter configuration from executing correctly.

PingFederate 7.2 - June 2014

- Multiple Virtual Server IDs
- OpenID Connect-based Centralized Session Management
- LDAP Enhancements
 - Improved performance on LDAP bind operations
 - Ability to control timeout on LDAP attribute lookups
 - Better connection cleanup when an idle LDAP connection is removed from the pool
- Administrative API Enhancements
 - IdP and SP Adapters
 - Access Token Management (OAuth)

- Custom HTTP Response Headers
- Improved Target Resource Validation
- Outbound Provisioning Monitoring
- New Username Token Processor
- Redesigned Default User-Facing Screens
- Other Enhancements
 - Support for Java 8
 - Support for OAuth Symmetric Proof of Possession for Code Extension
 - Upgraded JGroups to 3.4.4.Final
 - Various security enhancements
 - Over 50 other product issues resolved

PingFederate 7.1.4 - June 2014

- Addressed a potential CSRF issue with the PingFederate Administration console.
- Updated the default Multiple SSO in Progress template to escape additional input variables (`speed.bump.template.html`).
- Addressed an issue where Office 365 Outlook users can have their user accounts to be locked out after they change their Active Directory passwords.
- Corrected the Content-Type header for interoperability with Office 365 OneDrive.
- Made available the `$HttpServletRequest` object in the IdP logout template (`template.idp.logout.success.page.template.html`).
- Resolve an issue with body-less HTML Form POST messages sent by Internet Explorer when HTML Form based authentication follows an unsuccessful IWA authentication within a Composite Adapter.
- Made an improvement to correctly handle LDAP queries for Distinguished Names having forward slashes (/).

PingFederate 7.1 R3 - March 2014

- Administrative API Enhancements
 - Server Settings (OAuth Use Cases)
 - OAuth Settings
 - Authorization Server Settings
 - OpenID Connect Policy Management
 - Client Management
 - IdP Adapter Mapping
 - Access Token Mapping
- OAuth Enhancements
 - OAuth 2.0 Token Revocation (RFC7009); an index (`CLIENTIDIDX`) was added for the `client_id` column in the `pingfederate_access_grant` database table
 - Access Grant API
- RADIUS Support Extended
 - Challenge-Handshake Authentication Protocol (CHAP) Support
 - Vendor-Specific Attributes
 - NAS-IP-Address Passed in Access-Request
- Enhanced Support for Reverse Proxy Deployments
- Other Enhancements
 - Allow PingFederate's session cookie to be optionally configured to be persistent
 - Allow the Access Token Verification/Validation Grant Type to be used in combination with other grant types
 - Various security enhancements
 - Over 20 other product issues resolved

PingFederate 7.1.3 - February 2014

- Corrected potential security vulnerability.

PingFederate 7.1.2 - December 2013

- Corrected an issue with artifact binding in clustered deployments that resulted in sporadic failed transactions. This issue only applied to version 7.1.1.

PingFederate 7.1 R2 - December 2013

- Administrative API
- Tracking ID Enhancements
- Outbound Provisioning
 - Microsoft SQL Server Support added for Internal Provisioning Data Store
 - OAuth 2.0 Bearer Token Authorization for SCIM 1.1 Plugin
 - Ability to Define Deprovision Method for SCIM 1.1 Plugin
- OAuth and OpenID Connect Enhancements
 - Extended OAuth to allow multiple scopes to be grouped together and referenced as a single scope, enabling scope downgrade during token refresh
 - Allow administrators to define a maximum token lifetime for OAuth Access Tokens
 - OpenID Connect now supports requesting that the Authorization Server use specified authentication contexts when processing the authentication request
- Enhanced the PingFederate SP server to validate the Audience element in SAML assertions against the Entity ID defined in Server Settings and the Virtual Server IDs defined within the IdP connection
- Added TemplateRendererUtil class and template-render-adapter-example to PingFederate SDK
- Allow administrators to define a maximum token lifetime for OAuth Access Tokens
- Extended HTML Form Adapter to allow usernames to be stored and pre-populated in the login form after a user's first successful authentication
- OpenID Connect now supports requesting that the Authorization Server use specified authentication contexts when processing the authentication request
- Partner Entity ID (Connection ID) now available as input parameter to all Identity Store Provisioner requests
- Extended OAuth to allow multiple scopes to be grouped together and referenced as a single scope, allowing scope downgrade during token refresh
- Reintroduced Luna HSM Support
- Upgraded Jetty to 8.1.14
- Upgraded JGroups to 3.3.4.Final
- Various security enhancements
- Over 40 other product issues resolved

PingFederate 7.1.1 - November 2013

- Corrected potential security vulnerability
- Fixed XML namespace issue that resulted in failed SSO transactions at partners with hardcoded dependencies on the `samlp` prefix
- Resolved XML parsing issue that resulted in PingFederate rejecting incoming authentication requests
- Corrected logging issue that resulted in AJP-based transactions to fail

PingFederate 7.1 - August 2013

- Added support for Outbound Provisioning via PingOne
- Added identity store SDK for Inbound Provisioning
- Provided Remote Authentication Dial-In User Service (RADIUS) Support for console logon and password credential validation

- Added hierarchical (parent/child) configurations for plug-in instances
- Related plug-in instance override capability added for connections
- Enabled overriding the global Default Target URL for connections and adapter-to-adapter mapping
- Made STS client authentication details available for token authorization
- Extended JMX runtime monitoring support
- Enabled usage checking for certificates, adapters, token translators, credential validators and data stores
- Store OAuth client configuration as XML files (new default, database storage optional)
- Upgraded Jetty to 8.1.9
- Upgraded JGroups to 3.3.0.Final
- Various security enhancements
- Over 80 other product issues resolved

PingFederate 7.0.1 - May 2013

Addressed potential security vulnerabilities found since the PingFederate 7.0 initial, limited-availability release.

PingFederate 7.0 - April 2013

- Added support for System for Cross-domain Identity Management (SCIM)
 - Inbound Provisioning
 - Outbound Provisioning
- Initial Support for OpenID Connect
- New Context sources available in Token Authorization and Attribute Mapping workflows
 - HTTP Request
 - *Client IP* Request
- Administrative Console Enhancements
- Upgraded Jetty to 8.1.8
- Added support for the SAML AuthnRequest Scoping element Request (see *Administrator's Manual: Configuring the Requested AuthN Context Selector*)
- Created the Cluster Node Adapter Selector for use within Adapter Selection Request (see *Administrator's Manual: Configuring the Cluster Node Selector*)
- Allow PingFederate runtime context path to be customized Request
- Security Enhancements
- Over 50 product issues resolved

PingFederate 6.11 - December 2012

- Added password update via HTML Form Adapter
- Enhanced IdP Adapter Selector functionality:
 - Decision Tree Processing (see *Administrator's Manual: Mapping Selector Results to Adapter Instances*)
 - Connection Set Selector (see *Administrator's Manual: Configuring the Connection Set Selector*)
 - HTTP Header Selector (see *Administrator's Manual: Configuring the HTTP Header Selector*)
 - OAuth Scope Selector (see *Administrator's Manual: Configuring the OAuth Scope Selector*)
- Added localization for user-facing Web pages
- Added capability of defining globally an HTTP Header containing client IP addresses (see *Administrator's Manual: Defining an HTTP Header for Client IP Addresses*)
- Added OAuth fine-grain scope handling and JWT access-token support
- Added Dynamic JVM Resource Allocation
- Upgraded JGroups to 3.3.0.Alpha1 (snapshot from 10/30/12)
- Extended the Consent Form template to include the \$entityId parameter

- Extended the OAuth Grants Management template to show grant type, scope, and the date/time a grant was issued and updated
- Improved SLO handling in scenarios where a user performs multiple SSO requests
- Improved the ability for administrators to define validation rules for HTTP request parameters
- Improved JDBC data store attribute lookup functionality to allow for retrieval of multiple values from a single database column
- Made security enhancements
- Removed the requirement to include TokenProcessorId/TokenGeneratorId query parameters for STS requests to an IdP/SP connection when only a single instance of the token-processor/generator type is configured for the connection
- Enabled PingFederate to use Jetty's NIO (New I/O) backed SSL Connectors by default
- Enhanced Microsoft Office 365 support to allow for Kerberos interoperability with active clients built on top of Microsoft Online Services Sign-In Assistant
- Enhanced LDAP error code handling in the LDAP Username Password Credential Validator to enable more specific error messages to be provided to end users
- Enabled PingFederate to interoperate with Microsoft Dynamics CRM 2011

PingFederate 6.10.1 - January 2013

- Replaced OpenToken Adapter with version 2.5.1 to capture security enhancements
- Other changes to address potential security vulnerabilities

PingFederate 6.10 - September 2012

- Token Authorization
- STS token exchange mapping
- OAuth client mutual TLS authentication
- OAuth 2.0 final draft compliance

PingFederate 6.9 - June 2012

- Microsoft Office 365 interoperability
- STS transaction events logged to audit log
- Upgraded Jetty and removed underlying JBoss infrastructure

PingFederate 6.8 - April 2012

- Added centralized AD Domain/Kerberos Realm configuration
- Added OAuth Client Management REST API
- Added optional expiration of OAuth persistent grants
- Added multiple redirect URIs per OAuth client
- Added optional restricted scope subsets per OAuth client
- Added configurable consent page omission per OAuth client
- Added OAuth transaction events logged to audit log

PingFederate 6.7 - February 2012

- Added Splunk Application for PingFederate
- Improved administrative console navigation and save performance
- Added LDAP connection pooling options for LDAP data stores

PingFederate 6.6 - December 2011

- Added contextual IdP Adapter selection using Adapter Selectors
- Added ability to chain multiple IdP adapters together using the Composite Adapter

- Added the ability to use multiple IdP data stores for attribute retrieval and mapping into an IdP attribute contract
- Added an HTML Form Adapter and HTTP Basic Adapter to replace the LDAP Authentication adapter
- Support for the OAuth SAML 2.0 Bearer Assertion Grant Type
- Added an Admin Console Help system updater
- Added IPv6 support


PingFederate 6.5.2 - November 2011

- Security update since the PingFederate 6.5.1 release

PingFederate 6.5.1 - October 2011

- Security updates since the PingFederate 6.5 release

PingFederate 6.5 - August 2011

 **Note:** The PingFederate 6.5 release includes the features described below as well features that were added in a limited-distribution "Preview" release, described in the next section.

- PingFederate now functions as an OAuth 2.0 Authorization Server
- Added support for Thales (nCIPHER) nShield Connect HSM
- Account Linking can use an LDAP directory for a persistent data store in addition to a relational database system
- User-Defined Attribute Namespaces can be specified for Browser SSO protocols (similar to what was added to WS-Trust STS) to allow for better Microsoft interoperability
- Adapter to Adapter mapping now counts as a licensed connection
- LDAP Adapter updated to 2.2 with new default web form login template
- Jetty version upgrade from 6.1.7 to 6.1.26

PingFederate 6.5-Preview - April 2011

- Full STS metadata Claims Provider and Relying Party interoperability with Microsoft WIF, WCF, and ADFS 2.0
- Support multiple token-processor instances of the same token type
- Added SAML HoK subject confirmation in the SAML Token Generator
- Added option for STS SAML token KeyInfo to use a signing certificate reference rather than the full signing certificate
- Session-state modifications to support simultaneous and nested SSO transactions
- IdP adapter session handling for IdP adapters that rely on PingFederate for session management to allow for consecutive requests without prompting for credentials

PingFederate 6.4.1 - February 2011


- Corrected license expiration date calculation

PingFederate 6.4 - December 2010

- Support standard .NET WS-Trust Federation Bindings
- Support SAML 2.0 token Holder of Key (HoK) subject confirmation
- Added Metadata Exchange (MEX) endpoint for WIF client to generate bindings automatically for Username, X.509, and SAML tokens
- Added support for WS-Trust 1.4 ActAs property
- Added two-factor authentication capability with the VeriSign® Identity Protection (VIP) Authentication Service Adapter
- OpenToken Adapter 2.4.1 updated to correct issue with Cookie Transport Method and Replay Prevention
- Expanded digital signature secure hash algorithm types - SHA1, SHA256, SHA 384, and SHA512
- The provisioning log can be written to a database—Oracle, Microsoft SQL Server, and MySQL databases supported

- Added SAML protocol support for AuthnContextDeclRefs

PingFederate 6.3 - August 2010

 **Note:** The PingFederate 6.3 release includes the features described below as well features that were added in a limited-distribution "Preview" release, described in the next section.

- PingFederate STS claims-based identity capabilities extended to support interoperability with Microsoft WIF and WCF client frameworks
- Expanded SNMP monitoring variables available in the management information base (MIB)
- Increase the default PingFederate HTTP header buffer size to 8k
- Key stores and key store passwords are dynamically generated per installation
- The default SSL server certificate is generated upon initial startup if an SSL certificate does not exist
- LDAPS trust configuration no longer requires a server restart to take effect


PingFederate 6.3-Preview - April 2010

- Added support for logging to the ArcSight Common Event Format (CEF)
- Added ability to log to a database with failover to file—Oracle, SQL Server, and MySQL databases supported
- Added ability to disable automatic multi-connection validation if the validation time is causing excessive delay
- Extended JDBC Express Provisioning to support MS SQL Server stored procedures
- Added replay prevention capability to the OpenToken IdP Adapter bundled with PingFederate

PingFederate 6.2 - February 2010

- Added IdP-to-SP adapter mapping, which allows user attributes from an IdP adapter to be directly mapped to an SP adapter on the same PingFederate server to create an authenticated session or security context, without the need to generate SAML messages in between
- Provides enhanced logging capabilities including a new audit log, logfilter utility, and ability to log to any accessible file-server directory
- Provides enhanced support for configuration automation including certificate and key management, configuration archive management, and ancillary deployment files
- Extended JDBC Express Provisioning to support MS SQL Server Identity column types
- Added a Logout Endpoint to the LDAP Authentication Adapter
- In clustered mode, the default Inter-Request State Management methodology is now group RPC-based instead of cookie-based
- Added ability to extract CN from DN and extract username from email address for provisioner attributes
- In Luna HSM mode, added the ability to specify the location to store Trusted CA certificates, either in the Sun Java key store or the Luna HSM (see the configuration file `org.sourceid.config.CoreConfig.xml` in the `pingfederate/data/config-store` directory)

PingFederate 6.1 - September 2009

 **Note:** The PingFederate 6.1 release includes the features described below as well features that were added in a limited-distribution "Preview" release, described in the next section.

- Provides support for simplified PingFederate Express connection configuration and export
- Extends support for configuration automation, including listing, copying, and updating features for SaaS Provisioning channels and for Token Translators
- Provides enhanced support for SaaS Provisioning Health and Status Monitoring via JMX
- Provides licensing enhancements including support for organizational licenses, licenses that contain international characters, and web based license import
- Enhances the trust model to include support for anchored certificates, which allows certificates to be included in federation-transaction messaging and used for signature verification if, the given certificate matches the registered Subject DN and is issued by a certificate authority registered as a Trusted CA with PingFederate
- Supports "SP Lite", "IdP Lite", and "e-Gov" Liberty Interoperability profiles for SAML 2.0

PingFederate 6.1-Preview - June 2009

- Provides support for Express Provisioning to a JDBC data source
- Extends support for configuration automation, including listing, copy and update features for data stores and server settings
- Supports certificate-based authentication to the PingFederate administrative console
- Added UI-based data-archive deployment, as well as better error handling for common errors encountered
- Supports mapping of attributes passed in via a WS-Trust STS request to the outgoing token
- Supports logging of a tracking ID (transaction ID) associated with every log entry for a given request to the PingFederate server
- Corrects defects reported by customers in the previous release

PingFederate 6.0 - March 2009

- PingFederate now includes a WS-Trust Security Token Service (STS), enabling organizations to extend identity management to Web Services. The PingFederate STS shares the core functionality of PingFederate, including console administration, identity and attribute mapping, and certificate security management.
- PingFederate extends support for configuration automation, including connection management and adapter management via the existing command-line tool.
- PingFederate supports enhanced connection based licensing capabilities.
- PingFederate provides transaction based licensing capabilities for evaluation phase license enforcement.
- PingFederate allows administrators to specify the use of a separate certificate that is used for access to the administrative console and a different certificate for runtime processing.
- PingFederate supports configuration of LDAP Groups who are allowed access to the PingFederate Admin application based on PingFederate defined roles.
- PingFederate supports definition of LDAP data stores such that the connection URI for multiple LDAP servers can be specified as the connection string for that LDAP data store.
- PingFederate supports the Virtual List View (VLV) paging mechanism for retrieval of subsets of large result sets returned from the source LDAP data store during the provisioning process. This can significantly enhance performance for retrieval of data from Sun Directory Server (SDS) and similar LDAP servers that support VLV paging.
- PingFederate now stores the PingFederate software version within the configuration store.
- A number of defects reported by customers that existed in the previous release of PingFederate were addressed.

PingFederate 5.3 - December 2008

- PingFederate can be run as a service on Windows 64-bit platforms in addition to Windows 32-bit platforms and Linux platforms.
- PingFederate now supports deployment of SaaS Provisioning plug-ins (JAR files) via a separate installation package.
- PingFederate now supports automating configuration via a command line utility for connection management.
- PingFederate now supports capabilities for monitoring and control of the SaaS Provisioning configuration and data via a command line tool.
- PingFederate now supports validation of certificate revocation information via OCSP.
- PingFederate can now be deployed on Java 6 (JDK 1.6) platforms.
- PingFederate supports access to additional parameters on both the IdP side and the SP side via OGNL expressions.
- PingFederate supports "SP Lite" and "IdP Lite" Liberty Interoperability profiles for SAML 2.0.
- PingFederate supports configuration of the name, domain, and path for the cookie used for conveying state information between servers when cookie-based clustering has been configured.
- A number of past Known Issues and Limitations were addressed.

PingFederate 5.2 - August 2008

- A PingFederate IdP server now provides support for provisioning to selected SaaS providers. PingFederate supports provisioning of user account data from LDAP directories including Active Directory and Sun Directory Server. PingFederate stores synchronization data in JDBC data stores including Hypersonic (for demonstration purposes) and Oracle.
- PingFederate supports quick-connection templates to selected SaaS Providers, including Google Apps, and Salesforce.com

PingFederate 5.1.1 - July 2008

This release corrected several issues, including:

- SP signature verification was failing for assertions containing UTF-8 characters.
- In Windows the PingFederate server was unable to start when the JAVA_HOME system variable contained a space.
- Versions of the OpenToken library were placed in the wrong directory.
- Single Logout (SLO) with two SPs was not being performed for the IdP session(s).
- For SLO with three or more SPs, SP sessions were being stranded.
- Specific to PingFederate 5.1, Custom Data Sources no longer could be used for Adapter Contract fulfillment.
- When testing certain types of OGNL expressions, important error details were being lost when evaluation of these expressions failed.
- For SLO with at least two SPs, under certain circumstances error messages from SPs that did not initiate the SLO were not being processed correctly by the IdP.
- A PingFederate SP instance, when used with the OpenToken Adapter, was converting a plus "+" character to a space " " when constructing the URL for final redirect.
- Signature validation was failing within a PingFederate SP instance when it received an SLO message in which the SAML_SUBJECT was being encrypted.
- PingFederate 5.1 SP instance was no longer supported SiteMinder SSO Zones.

PingFederate 5.1 - April 2008

- The default behavior when PingFederate cannot access a Certificate Revocation List (CRL) is now set correctly. The server no longer treats a non-retrievable CRL as a reason to label certificates as revoked. CRL processing behavior is managed by the revocation-checking-config.xml file in the .
- The IP address to which PingFederate's SNMP agent binds is now controlled by the pf.monitor.bind.address property in the run.properties file.
- Building either of the two example adapters included in the PingFederate SDK no longer fails with an error regarding a missing README.txt.
- The PingFederate server now correctly maintains temporary files within the . The server no longer writes temporary files to the tmp directory of the user running the server.
- Express Provisioning allows user accounts to be created in an LDAP repository and updated directly by an SP PingFederate. User provisioning occurs as part of SSO processing and may be used with any IdP partner.
- The Signature Policy screen in SAML IdP connections contains improved language clarifying how signatures are used to guarantee authenticity of SAML messages.
- The PingFederate package now contains v6.1.7 of Jetty. Jetty is the servlet container used by PingFederate.
- The PingFederate SDK contains a ConfigurationListener interface that may be utilized by developers building adapters. This interface contains methods invoked by the server in response to certain adapter-instance lifecycle events such as creation and deletion.
- Adapter-instance Summary screens now display adapter-instance configuration values specified within a TableDescriptor.
- After the third consecutive failed login attempt, an administrator is blocked from accessing the administrative console for a configurable amount of time (default = 60 seconds).
- When changing an administrator password, the server now forces the new password to differ from the existing password.

- Access to services exposed by the PingFederate server now requires client authentication. These services include Attribute Query, JMX, and Connection Management. An administrator may choose to require client authentication for access to the SSO Directory Service. An ID and Shared Secret comprise the credentials needed for authentication.
- For security, the use of "Expression" in contract fulfillment screens is now disallowed by default. For backward compatibility, customers deploying a configuration archive from a previous version of PingFederate in which expressions were used will continue to have access to expressions. Allowing expressions creates a potential security concern in the PingFederate administrative console.
- The Quick-Start SP Application no longer uses an OGNL expression in fulfilling the SP adapter contract.
- HTTP TRACE requests sent to PingFederate now result in an HTTP 403 Forbidden response.
- By default, "weak" ciphers are no longer supported during SSL handshaking. (For more information as to which cipher suites the server supports, examine the `com.pingidentity.crypto.SunJCEManager.xml` file.)
- It is no longer possible for an administrator to circumvent role and access permissions within the administrative console by direct URL access. The server evaluates HTTP requests for a URL against an administrator's assigned role(s) and responds appropriately.
- The PingFederate runtime server's HTTP listener is now turned off by default. Only messages sent over HTTPS are accepted. This may be controlled in the `run.properties` file.
- Use of class and package names specific to a PingFederate version were removed from sample source code contained in the PingFederate SDK.
- The `/idp/startSSO.ping` endpoint now supports an optional ACSIdx query parameter for SAML v2 partners. When provided, the PingFederate IdP attempts to send the SAML Assertion to the Assertion Consumer Service corresponding to the specified Index.
- The initial and maximum JVM heap sizes are set to 256 MB and 1024 MB, respectively, by default. These changes should improve runtime performance on servers with sufficient memory. These settings reside in the `run.bat` and `run.sh` files of the .
- During server startup, PingFederate now reports relevant environment variables and adapter- instance information to the `server.log`.
- Existing partner connections can be deleted through a SOAP call from an external client application.
- The `pf-legacy-runtime.war` file is no longer deployed by default. This WAR file allows a PingFederate server to continue support of legacy endpoints (those endpoints supported by PingFederate 2). When replacing an existing PingFederate 2 server deployment, manually move this WAR to the .
- The PingFederate server can be configured to support the use of a proxy server when retrieving a CRL from a Certificate Authority. Relevant configuration settings reside in .
- When the PingFederate server relies upon an external LDAP directory to authenticate administrative users, the `ldap.password` property in the `pingfederate/bin/ldap.properties` file now supports encrypted credentials.
- The PingFederate server allows imported SSL server certificates containing signatures from one or more intermediate Certificate Authorities. When SSL clients request an SSL connection to the PingFederate server, the entire SSL server certificate chain is presented.
- The Summary and Activation screens for both IdP and SP connections display a valid URL that serves as an example of a `startSSO.ping` endpoint used by local applications integrating with PingFederate.
- The Web SSO entry screens for both IdP and SP connections include summary information in a table describing relevant configuration settings.
- The PingFederate server prevents auditors from accessing links on the Main Menu that impact external resources. This includes exporting SAML metadata, signing XML files, and creating configuration archives.

PingFederate 5.0.2 - March 2008

- IdP Persistent Reference Cookie (IPRC) — Provides a mechanism allowing an SP PingFederate server to discover a user's IdP based on a persistent browser cookie that contains a reference to the IdP partner previously used for SSO. This feature provides an alternative to standard IdP Discovery for SP-initiated SSO, as defined in the SAML specifications, which uses a common-domain cookie (CDC) written by the IdP (see the *PingFederate Administrator's Manual*). Unlike the IdP Discovery cookie, the IPRC is written by the SP PingFederate each time an SSO event for the user occurs (either IdP- or SP-initiated). The cookie identifies the IdP partner using information in the SAML assertion. For subsequent SP-initiated SSO requests, the SP server can skip a previously

required step prompting the user to select an IdP for authentication when multiple IdP partners are configured but none is specifically identified in the SSO call received by the SP PingFederate server.

- Updated the IdP-selection template to make it easier to use. The new selection template is used when no IPRC (or CDC) is available and when there are multiple IdPs to which the user might have previously authenticated.
- Corrected an issue in which a Concurrent Modification Exception was encountered when server clustering is used and debug is turned on for log files. (The workaround for this issue in previous releases is to turn debug off.)
- When no certificate revocation list (CRL) is found during certificate validation checking, the subject certificate is assumed to be valid for the current SSO/SLO transaction. Previously, when no CRL was found, the certificate was deemed invalid and the transaction aborted. The default setting is changed for this release to prevent problems with upgrading to PingFederate 5.x from previous versions.

PingFederate 5.0.1 - January 2008

- Support for rapid provisioning of partner connections is available using Auto-Connect technology. Leveraging the existing SAML 2.0 specification, Auto-Connect allows PingFederate deployments to scale easily with minimal manual involvement. The majority of partner connection configuration occurs at runtime through the exchange of dynamically generated metadata.
- Administrators can authenticate to the administrative console using credentials in an external LDAP directory. This allows organizations with existing admin accounts to provide access to the console without creating and managing individual accounts within PingFederate.
- Partner connections may be created by importing them programmatically into PingFederate through a SOAP interface. This allows administrators to provision partner connections without accessing the administrative console manually. The Connection Management screens (both IdP and SP) contain an "Export" action that creates an XML file containing a connection's configuration.
- Server configuration data may be replicated to a cluster through a SOAP call to the administrative console. This allows cluster deployments to receive configuration changes without accessing the administrative console manually.
- The SAML 2 Attribute Query Profile is supported by PingFederate. This allows SPs to request user attributes from an IdP independent of user authentication.
- Multi-valued attributes passed in an Assertion to a WS-Federation partner conform to how ADFS expects them.
- SAML metadata may be generated with a digital signature to guarantee authenticity.
- The Protocol Endpoints popup contains online help links. These links may be used to learn more about the server's endpoints from a partner perspective.
- The User-Session Creation screen in the IdP connection flow contains summary information that provides administrators with insight into the current configuration. Similarly, the Assertion Creation screen in the SP connection flow also provides administrators with useful configuration information.
- Administrators can specify a descriptive "Connection Name" for partner connections.
- The "Web SSO" portion of partner configuration is distinct from first/last-mile configuration.
- Connection summary screens contain +/- buttons to show/hide major sections.
- Metadata import extracts the Base URL from the metadata file and populates relative URLs within the connection.
- PingFederate communicates with an SNMP network-management console using standard SNMP Get and Trap operations.
- The PingFederate engine exposes a URL designed for load balancers to determine whether a PingFederate server is available to process transactions.
- Certificate-management usability is improved.
- Certificate expirations are tracked by PingFederate, and impending expirations may result in a notification sent via email, when configured.
- New, more user-friendly sample applications focus on demonstrating PingFederate server functionality (Quick-Start Guide).
- The entry screen into the "Credentials" area of connections contains useful information about the credentials used.
- The server ID is no longer displayed to the administrator.
- A cluster's administrative console no longer aggregates transactions counts.
- The "Cluster Management" link replaces "High Availability" on the Main Menu.

- Clustering supports TCP and UDP, node authentication, optional encryption, and use of a single port for all communication.
- CRL processing updated to support the U.S. GSA's E-Authentication v2 specification.
- The "About" pop-up contains additional license-key information.

PingFederate 4.4.2 - October 2007

- Mitigated a number of potential security vulnerabilities regarding XML document processing

PingFederate 4.4.1 - June 2007

- Addresses user-interface defects related to attribute query, XML encryption, and LDAP data store lookups

PingFederate 4.4 - May 2007

- Removal of support for the U.S. GSA's E-Authentication v1.0 specification
- Addition for support of signed metadata files
- Support for partner certificate revocation through CRLs
- Increased flexibility around encryption of Name ID in SAML v2.0 SLO requests when the Name ID is encrypted within an assertion
- Support for the SOAP binding for both inbound and outbound SAML v2.0 messages
- Improved support for deployments where the server contains multiple network interfaces
- Usability enhancements to the Main Menu layout and Local Settings flow
- More sophisticated attribute-fulfillment operations through support of a Java-like syntax for data manipulation
- Removal of extraneous credentials settings for WS-Federation and SAML v1.x connections
- Improved display of long connection IDs on the Main Menu
- Inclusion of a demo application that complements the existing sample applications as described in the Quick-Start Guide

PingFederate 4.3 - March 2007

- Virtual Server Identities allow PingFederate to use distinct protocol identifiers in the context of a particular partner connection
- Additional customizable end-user error pages for 'page expired' and general unexpected error conditions
- Increased flexibility by allowing for a list of additional valid hostnames to be used for incoming protocol message validation
- Optionally, the SSO Directory Web Services can be protected with HTTP basic authentication
- New administrative console error page
- Improved short-term state management memory utilization for improved system resiliency
- Improved input-data validation and character-entity encoding of data when displayed--for protection against cross-site scripting attacks
- An IdP connection configured to use only a single SP Adapter Instance will ignore the URL-to-Adapter mapping step at runtime and just use the given adapter
- Blocked directory indexing to limit browsing of static web content
- Disabled unnecessary JRMP JMX port usage
- Mitigated HTTP response splitting attacks by disallowing potentially dangerous characters in all redirects

PingFederate 4.2 - December 2006

- Enhanced transaction logging functionality
- Sensitive user attribute values can be masked in log files to enhance privacy considerations
- The administrative console runs on a distinct port from the runtime engine allowing for more flexible and secure deployment options

- New filtering functionality on connection management screens enables easier management of large numbers of federation partners
- Adapter SDK enhancements to facilitate file downloads
- Usability refinements on X.509 certificate summary screens
- Less verbose description of certificates in drop down boxes improve look and feel
- Multiple partner endpoints of the same type can be configured to use the same binding
- Improved support for reverse proxy deployments

PingFederate 4.1 - October 2006

- Liberty Alliance interoperability certified
- SAML2 x509 Attribute Sharing Profile (XASP)
- Optional Hardware Security Module (HSM) mode, that enables storage of private keys and crypto processing on an external HSM unit that is FIPS-140-2 certified
- Updated Protocol Configuration Wizard
- Error handling templates that can be used to build SSO/SLO landing pages that communicate error status and support instructions to users
- Configuration options that enable multiple, simultaneous authentication profiles for the SOAP back-channel, including HTTP Basic, SSL Client Certificates, and Digital Signatures
- Digital signature capability for client authentication when using SAML 1.x
- Pop-up server endpoint display that filters by role and configurations made
- Two digital signature verification certificates can be assigned to a connection, allowing the partner flexibility in selecting one certificate or the other. When one certificate expires, the other certificate is used without the need for close synchronization.
- A `run.properties` configuration that allows an admin to specify an alternate port with which to communicate over the back-channel to partner's SAML gateway
- Support for 32-and 64- bit machine architectures

PingFederate 4.0 - June 2006

- Deploy multiple adapters as an IdP to look up different session security contexts across security domains and applications
- Save a partially completed connection as a draft
- Copy a connection to rapidly set up other partners or test environments with similar configurations
- Attribute source SDK enables retrieval of attributes from additional data source interfaces such as SOAP, flat files, or custom interfaces
- Multi-administrator support
- Ability to edit SP adapters that are in-use with target systems
- Encrypt or decrypt entire assertions or select elements, of particular value when intermediaries may handle SAML traffic
- Generate unique, Transient Name Ids each time the user federates to protect their identity
- SAML 2-compliant IdP Discovery mechanism that enables an SP to dynamically determine the appropriate IdP for the user
- Integration Kits provide additional methods that streamline passing of authentication context from an IdP to an SP
- Single log-out across all connections and protocols that support SLO
- Using an affiliate id, an SP can instruct an IdP to re-use the same persistent name identifier that was already used at other applications within the portal
- Non-normative support for SP-initiated SSO with SAML 1.x protocols

PingFederate 3.0.2 - February 2006

- Upgrade of Jetty component to v5.1.10 in response to a security warning from the National Vulnerability Database

PingFederate 3.0.1 - December 2005

- Complete clustering support
- Optional email notification on licensing issues
- You can edit a previously configured connection (either IdP or SP) that uses a data store that is unavailable

PingFederate 3.0 - November 2005

- Support for SAML 2.0
- Use-case wizard for partner connection configurations
- Support for multiple security domains
- Redesigned user interface
- Embedded clustering
- Fixes for LDAP and JDBC connectivity

PingFederate 2.1 - July 2005


- Patched a concurrency bug in the XML security library
- Patched a memory leak in the XML-to-object binding library
- Removed the core protocol processor's reliance on a workflow engine to resolve a memory leak and improve overall performance
- Fixed a subtle memory leak in the module that tracks assertions in order to prevent replay in the POST profile
- Updated the default server SSL certificate (extended the expiration date)

PingFederate 2.0 - February 2005

- Initial release

PingFederate Upgrade Utility user guide

This Upgrade Utility is a command-line tool that automatically migrates existing PingFederate® 6.x (or more recent) configurations to the current version. The tool first installs the new distribution and then copies over files or property values from the previous installation.

 **Note:** The Upgrade Utility does *not* affect the existing installation of PingFederate.

As an option, you can run the tool with a *custom* switch that allows you to choose whether to override any new system defaults or to install newer versions of bundled adapters (see [Custom mode](#) on page 629).

The Upgrade Utility creates a detailed log showing what files and settings are affected during the upgrade.

If you are upgrading from PingFederate 5.x, contact support@pingidentity.com for more information.

Intended Audience


The PingFederate Upgrade Utility and this document are intended for system administrators with significant knowledge of the current PingFederate configuration and enterprise deployment.

Upgrading PingFederate

Fully upgrading your existing PingFederate installation involves initial preparation before running the Upgrade Utility. Then afterward, some manual copying may be needed for customized files or supporting Web-application deployments.


We recommend that you upgrade your test environment and verify that the new installation meets your expectations before upgrading your production environment. After you complete the upgrade process successfully, you may create a backup of your previous installation and remove it from your server.

The Upgrade Utility is located in the `/bin` directory of the distribution ZIP as either a batch file (for Windows installations) or a shell script (for UNIX/Linux installations).

 **Note:** The Upgrade Utility does not migrate OAuth clients that are created from PingFederate 6.5 through 6.7 or stored in the internal HSQLDB (formerly known as HyperSQL and Hypersonic) database since PingFederate 6.8. Use any of the following interfaces to reconfigure your OAuth clients:

- The **OAuth Settings > Client Management** screen 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 higher) to store OAuth clients in an external database, the new installation retains that configuration as well.

 **Important:** Specific modifications made since PingFederate 6.11, such as security updates related to XML encryption, changes in LDAP filter, and updates/upgrades of third-party library files, may affect existing deployments. For detailed information, review [Upgrade considerations](#) on page 578 in the PingFederate *Release notes*.

Integration Kits

The Upgrade Utility copies Adapters, Connectors, Token Translators, and Integration Kits, from the previous installation to the new installation; the tool does not upgrade them automatically. You may download newer versions by visiting the Ping Identity [Downloads](#) website and upgrade the kits manually.



Note: Documentation for Integration Kits are available from the Ping Identity [documentation](#) website.

PingFederate installer

Depending on the source installation, you may also upgrade PingFederate by using the platform-specific installer. The conditions are described as follows:

Source installation	Possible upgrade paths on a Windows server	Possible upgrade paths on a Red Hat Enterprise Linux server	Upgrade path on a UNIX/Linux server
PingFederate 6.x to 8.x installed by using the PingFederate product ZIP file	PingFederate Upgrade Utility	PingFederate install script PingFederate Upgrade Utility	PingFederate Upgrade Utility
PingFederate 8.x installed by using the PingFederate platform-specific installer	PingFederate installer for Windows PingFederate Upgrade Utility	PingFederate install script PingFederate Upgrade Utility	PingFederate Upgrade Utility

You may download the PingFederate platform-specific installers from the [PingFederate Downloads](#) website.

Prepare for the upgrade

Before running the Upgrade Utility, complete the following tasks:

- If you have not already done so, download the current version of the PingFederate ZIP file and obtain a new license key.
- If you have not recently updated the Oracle Java SE Runtime Environment (Server JRE) to version 8 on your PingFederate host machine(s), you must do so.

When you upgrade the Server JRE, be sure to 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 Java 8. 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.

- Optional: Complete any unfinished connections (*Drafts*) in the existing PingFederate administrative console, *if* you want to include them in the migration.
- Unzip the Upgrade Utility distribution into a directory on the same machine running PingFederate.
- If you are upgrading PingFederate in a clustered-server environment, unzip the Upgrade Utility on all the PingFederate servers.

Run the PingFederate Upgrade Utility

Follow these steps to upgrade your PingFederate server using the PingFederate Upgrade Utility.

If you are upgrading multiple PingFederate servers in a cluster, start with the console node.

1. Download the latest version of PingFederate and PingFederate Upgrade Utility from [PingFederate Downloads](#) website.
2. Unzip the Upgrade utility ZIP file into a directory on the same machine running PingFederate.

3. Stop PingFederate.
4. On the command line, change the current directory to the `bin` directory of the Upgrade Utility and execute the following command:

Operating system	Command
Windows	<code>upgrade <pf_install_source> <outputDir> <pingfederateZip> [-c]</code>
UNIX/Linux	<code>./upgrade.sh <pf_install_source> <outputDir> <pingfederateZip> [-c]</code>

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.

<outputDir>

The full or relative path of the directory that will contain the new PingFederate base directory

<pingfederateZip>

The full or relative path and file name for the current distribution

-c

The optional parameter to run the tool in custom mode, which allows administrators to override newer default security settings (if any) and to upgrade the program files for the OpenToken Adapter (see [Custom mode](#) on page 629).

After a moment, the command window displays messages indicating upgrade progress. The process is complete when this message appears:

```
Upgrade completed with [N] errors and [N] warnings
```

If there are errors, scroll up the command window to see them and correct indicated problems. Errors during the upgrade should be rare but may include such input/output problems as missing or malformed configuration files in the source installation.

The messages, including any errors, are also logged to `upgrade.log` in the Upgrade Utility base directory.



Tip: Rerun the Upgrade Utility as many times as needed to correct any problems.

5. Start the new PingFederate installation and open the administrative console.
6. Follow the on-screen instructions to upload your license file for the new version.
7. Verify configuration settings in the administrative console.



Note: Access the connection management screens and ensure that automatic validation is run on each connection.

8. If you are upgrading a clustered PingFederate environment:
 - a) Stop each engine node and run the Upgrade Utility to upgrade its PingFederate installation.
 - b) Start the *new* installation on each node.
 - c) Log on to the PingFederate administrative console.
 - d) Go to the **Server Configuration > Cluster Management** screen.
 - e) On the **Cluster Management** screen, ensure all nodes are shown in the list of IP addresses.
 - f) Click **Replicate Configuration**.
 - g) Restart PingFederate on each engine node.

9. If PingFederate is running as a service, re-install it (for Windows) or reconfigure it (for UNIX/Linux) to use the new PingFederate installation.

Operating system	Instruction
Windows	<ol style="list-style-type: none"> 1. Remove the existing PingFederate service (see Remove PingFederate from a Windows server). 2. Install the new PingFederate service (see Install the PingFederate service on Windows manually).
UNIX/Linux	<p>Modify the existing symlink to point to the new PingFederate installation.</p> <p>If the older version of PingFederate was installed using the PingFederate install script for Red Hat Enterprise Linux, update the existing symlink so it points to the new PingFederate installation. By default, the symlinks created by the install script are located in the <code>/usr/local/</code> directory and are named <code>pingfederate-<i><index></i></code>, where <i><index></i> is a value assigned by the installation program.</p>

For a clustered PingFederate environment, re-install or reconfigure the PingFederate service on all nodes.

Custom mode

The custom-mode feature (invoked with the `-c` option on the command line) allows administrators to override several newer default security settings (new, depending on which PingFederate® version is currently running). 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, if applicable.

Security defaults

In general, using the new security defaults is highly recommended and should not cause significant issues for most PingFederate installations. The newer 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. If you want to see which cipher suites are commented out, choose yes (y) when prompted on whether to use the new defaults. Then, after the upgrade is complete, refer to either of the following configuration files in the new installation's `<pf_install>/pingfederate/server/default/data/config-store` directory:
 - `com.pingidentity.crypto.SunJCEManager.xml`
 - `com.pingidentity.crypto.LunaJCEManager.xml`
- Disabling SQLFilter as of PingFederate version 6.2.
- Disabling the sending of PingFederate session cookies by way of the unsecure hypertext transfer protocol. As of version 6.5, the session cookie is configured to be sent only using HTTPS.

Adapter upgrade

Upgrading the OpenToken Adapter from previous versions is also recommended and will not normally require any follow-on configuration changes.

- If your existing installation uses a version of the OpenToken Adapter prior to 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 prior to 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.

Review post-upgrade tasks

Review the following post-upgrade tasks:

- [Copy customized files or settings](#) on page 630
- [Review database changes](#) on page 631
- [Review log configuration](#) on page 633
- [Migrate to the integrated Username Token Processor](#) on page 634

It is also important to perform runtime tests to ensure the new PingFederate® installation fulfills your existing use cases.

Copy customized files or settings

User-facing screens

If you have modified any Velocity templates for user-facing screens and you wish to preserve the customized user experience, you must copy the HTML files over 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.

Example: `<link rel="..." href="assets/css/screen.1.css"/>`

Likewise, add `.1` to any references in the copied templates to the installed image files contained in the `assets/images` directory.

Example: ``

Email notifications

Similarly, if you have modified the email notification templates prior to version 8.2, you must migrate custom changes manually to the new HTML-based templates to preserve those changes. The plain text templates (`message-template-*.txt`) can be found in the `<pf_install>/pingfederate/server/default/conf` directory in your previous installation. The new HTML-based templates are located in the `<pf_install>/pingfederate/server/default/conf/template/mail-notifications` directory 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 corresponding changes manually in the new PingFederate deployment.

For upgrading PingFederate version 6.9 (or higher), you can simply copy over the relevant files from the Jetty configuration directory, `<pf_install>/pingfederate/etc`. For versions prior to 6.9, first identify any changes made to the JBoss configuration, then make corresponding changes for the newer Jetty configuration.

For example, a commonly modified file prior to version 6.9 might be the `jboss-service.xml` file located in the `<pf_install>/pingfederate/server/default/deploy/jetty.sar/META-INF` directory. When changes are identified, 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`.

Other configuration files

Other configuration files stored in the `<pf_install>/pingfederate/server/default/conf` directory, such as the `cluster-*.conf` files for advanced clustering deployment architecture, are not migrated automatically. To preserve your custom settings (if any), create a backup of the current configuration files and merge your changes to the current files.

size-limits.conf

Prior to version 8.4.2, the `InterReqStateMgmtMapImpl.expiry.mins` setting in the `<pf_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
<code>InterReqStateMgmtMapImpl.expiry.mins.state.map</code>	Inter-request state information	30
<code>InterReqStateMgmtMapImpl.expiry.mins.attr.map</code>	Adapter session-state data	1440 (24 hours)

The new settings help 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.

Review 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. If your PingFederate environment connects to one or more database servers, review the following information and make changes accordingly.

An improved index in the database table for OAuth clients

Applicable only to customers who had configured PingFederate to store OAuth clients on a database server.

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.

You must modify the index in to 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. Consult with your database administrator, as needed.

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 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 4.0 (or older) drivers, contact your vendors for the latest drivers and replace the older JDBC database-driver JAR files with the latest.


Alternatively, if you prefer to upgrade the older JDBC drivers at a later time, use the following steps to revert the JDBC connection pool library to BoneCP:

1. Edit `<pf_install>/pingfederate/server/default/data/config-store/com.pingidentity.jdbc.DataSourceDeployerFactory.xml`.
2. Replace `dbcp` with `bonecp`.
3. Save the file.
4. Restart PingFederate.

If you have a clustered PingFederate environment:

- a. Perform the steps above on the console node.
- b. Sign on to the PingFederate administrative console.
- c. Go to the **Server Configuration > Cluster Management** screen.
- d. Click **Replicate Configuration** to push this change to the engine nodes.

 **Important:** Support for BoneCP as the JDBC connection pool library has been deprecated and will be removed in a future release.

 **Note:** PingFederate has been tested with vendor-specific JDBC 4.1 drivers. To obtain your database driver JAR file, contact your database vendor. Database driver file should be installed to the `<pf_install>/pingfederate/server/default/lib` directory. You must restart the server after installing the driver.

Changes in the database tables for log messages

Applicable only to customers who have configured Log4j or Log4j 2 to write log messages to database tables on Microsoft SQL Server.

For performance reasons, PingFederate 8.0 started using the NVARCHAR data type instead of the VARCHAR data type. 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

Applicable only to customers who have created a database table for account linking on Microsoft SQL Server.

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. The table-setup script, `account-linking-sqlserver.sql`, in the `<pf_install>/pingfederate/server/default/conf/account-linking/sql-scripts` directory 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

Applicable only to customers who have deployed Microsoft SQL Server to store their 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. The table-setup script, `oauth-client-management-sqlserver`, in the `<pf_install>/pingfederate/server/default/conf/oauth-client-management/sql-scripts` directory 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

Applicable only to customers who have deployed Microsoft SQL Server to host their OAuth grant data store.

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. The table-setup scripts (targeted for Microsoft SQL Server) in the `<pf_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's contract, which requires an update in the OAuth grant data store. If you are upgrading from PingFederate 7.2.1 or an earlier version with OAuth use cases, run the table-setup script, `access-grant-attribute-<database>.sql`, found in the `<pf_install>/pingfederate/server/default/conf/access-grant/sql-scripts` directory for your database server. Consult with your database administrator, as needed.

New indexes in the database table for OAuth persistent grants

Since PingFederate 7.1 R3, a couple of indexes have been added to the `pingfederate_access_grant` table.

PingFederate version	Column	Index
7.3	<code>expires</code>	<code>EXPIRESIDX</code>
7.1 R3	<code>client_id</code>	<code>CLIENTIDIDX</code>

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, `access-grant-<databaseServer>.sql`, found in the `<pf_install>/pingfederate/server/default/conf/access-grant/sql-scripts` directory. Consult with your database administrator, as needed.

Changes in a database table supporting nested group membership


Outbound provisioning of groups and nested group membership requires an update in the internal data store.

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, `add-subgroup-membership-<database>.sql`, found in the `<pf_install>/pingfederate/server/default/conf/provisioner/sql-scripts/updates` directory for your database server to update the existing `group_membership` table. Consult with your database administrator, as needed.

If you are upgrading from PingFederate 7.0.1 or an earlier version, consider using the table-setup script, `provisioner-<database>.sql`, found in the `<pf_install>/pingfederate/server/default/conf/provisioner/sql-scripts` directory for your database server as a reference to add the new `group_membership` table to your existing internal data store.

Alternatively, you may create a new database using the table-setup script and let the outbound provisioner repopulate the new internal data store on its next run, regardless of the current version of your PingFederate.

 **Caution:** The table-setup script removes the existing (outbound provisioning) tables. If you wish 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.

Consult with your database administrator, as needed.

Review log configuration

Starting with version 8.0, PingFederate® uses Log4j 2 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 the older version, PingFederate 8.0 starts using Log4j2 after the upgrade process.

If `log4j.xml` has changed since it was originally installed, the Upgrade Utility makes a backup of the customized `log4j.xml` file and creates a new `log4j2.xml` based on default settings. The backup is located in the `<pf_install>/pingfederate/server/default/conf` directory. The filename is `log4j-old-`

<SourceVersion>.xml, where <SourceVersion> is the version number of the source PingFederate installation.

Upgrading from PingFederate 8.x

If the Upgrade Utility determines that the Log4j2 configuration file (log4j2.xml in the <pf_install>/pingfederate/server/default/conf directory) has changed since it was originally installed, new features are *not* activated; the tool does not currently support automatic merging of customizations made to the existing logging configuration. Instead, the Upgrade Utility copies the modified log4j2.xml to the new installation intact and renames 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:

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

PingFederate® 6.x and 7.x use Log4j to write log messages.

If the Upgrade Utility determines that the Log4j configuration file (log4j.xml in the <pf_install>/pingfederate/server/default/conf directory) has changed since it was originally installed, the tool copies 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 Log4j 2:

1. Review custom changes by comparing log4j-old-<SourceVersion>.xml against log4j.xml that is located in the pf-upgrade-<NewVersion>/reference-files/<SourceVersion>/server/default/conf directory from the Upgrade Utility.



Important: The <SourceVersion> values must match each other.

2. Modify log4j2.xml to suit your needs.



Tip: The configuration syntax between Log4j and Log4j 2 varies. For more information, consult [Log4j 2 documentation](http://logging.apache.org/log4j/2.x/manual/index.html) (logging.apache.org/log4j/2.x/manual/index.html).

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.

Migrate to the integrated Username Token Processor

As of PingFederate® 7.2, the Username Token Translator (WS-Trust STS use cases) 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.


1. Create an instance of the integrated Username Token Processor.



Tip: In the **Type** screen, select **Username Token Processor** from the list.




Note: 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.

2. Map the new token processor instance to the applicable WS-Trust STS SP connection in the **IdP Token Processor Mapping** screen.
 -  **Note:** Repeat this step if you have multiple WS-Trust STS SP connections.
3. Test your WS-Trust STS use cases.
4. 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.
5. Delete all token processor instances of the deprecated Username Token Translator in the **IdP Configuration > Token Processors** screen.
6. Remove the `pf-username-token-translator-<version>.jar` file from the `<pf_install>/pingfederate/server/default/deploy` directory on all PingFederate servers.
7. Restart PingFederate on all PingFederate servers.

Migrate to the integrated LDAP PCV

As of PingFederate® 7.3, the integrated LDAP Username Password Credential Validator (PCV) is capable of returning 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**, it is recommended to migrate to the integrated LDAP PCV.

1. Create an instance of the integrated LDAP PCV.
 - a) On the **Server Configuration > Password Credential Validators** screen, click **Create New Instance**.
 - b) On the **Type** screen, enter the required information and select **LDAP Username Password Credential Validator** from the list.
 - c) On the **Instance Configuration** screen, select an LDAP data store from the list, enter a search base and a search filter, and select the scope of the search.
 -  **Tip:** You may reuse the information from the existing **LDAP PCV with Extended Attributes** instance.
 - d) On the **Extended Contract** screen, enter `memberOf` under **Extend the Contract** and click **Add**.
Other commonly used attributes are `DN`, `givenName`, `mail`, and `username`.
 - e) On the **Summary** screen, review the setup and click **Done**.
 - f) On the **Manage Credential Validator Instances** screen, click **Save**.
2. In configuration where the **LDAP PCV with Extended Attributes** instance is used, replace it with the newly created **LDAP Username Password Credential Validator** instance.

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. When the **LDAP PCV with Extended Attributes** instance is no longer in-use, delete it from the **Server Configuration > Password Credential Validators** screen.
4. Remove the `LDAPExtendedAttributesPCV-<version>.jar` file from the `<pf_install>/pingfederate/server/default/deploy` directory on all PingFederate servers.
5. Restart PingFederate on all PingFederate servers.

Verify the new installation

Integration kits and custom solutions may introduce third-party dependencies that could trigger runtime errors. We recommend that you perform runtime tests, verifying that the previously deployed use cases are functional in the new installation.

Known limitations

- Neither the Upgrade Utility nor the platform-specific installers migrates OAuth clients that are created from PingFederate® 6.5 through 6.7 or stored in the internal HSQLDB (formerly known as HyperSQL and Hypersonic) database since PingFederate 6.8. Use any of the following interfaces to reconfigure your OAuth clients:
 - The **OAuth Settings > Client Management** screen 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 higher) to store OAuth clients in an external database, the new installation retains that configuration as well.


- As of version 6.5, PingFederate supports Thales nShield Connect HSM configurations. If your PingFederate installation is configured in a clustered environment with 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.
- The OGNL library used by PingFederate was upgraded with version 6.4. If you use OGNL expressions, we recommend retesting the expressions using the PingFederate administrative console.
- As of PingFederate 6.2, TLS renegotiation was disabled by default for all HTTPS listeners/connectors. This security update affects only existing partner connections using WS-Trust STS and those making use of inbound mutual SSL/TLS authentication for the SOAP or artifact SAML bindings. In such cases, administrators must manually reconfigure `run.properties` to enable the secondary HTTPS port. In addition, affected connection partners must update their configurations to use the new endpoint.
- The Upgrade Utility does not migrate any data maintained in the PingFederate internal HSQLDB database or any external database. If, for example, SaaS Provisioning is enabled in the new PingFederate instance using the internal database, it is re-initialized from the provisioning source.

PingFederate performance tuning

The default configuration for PingFederate® 8.x is acceptable for most small size deployments; however, additional tuning is recommended for mission-critical and high-transaction volume deployments. Fine-tuning a few simple application and system level settings enables PingFederate to achieve maximum performance out of the hardware chosen for your deployment.

This topic addresses several areas of tuning such as logging, concurrency, memory, and Java-specific tuning options. The topic 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.

 **Tip:** The performance tuning guides for PingFederate 7.x and 6.x are available from support.pingidentity.com.

 **Note:** An additional document related to performance, the *PingFederate Capacity Planning Guide*, is also available to customers as a performance data reference. This document is available from support.pingidentity.com.


Logging

Logging is a feature for tracking various aspects of the system's operation. It requires a certain amount of system resources, which affects the system's overall performance. Of particular expense is the actual writing to the log files. PingFederate® 8.x uses the high-performance asynchronous logger from Log4j 2 to minimize the performance impact of logging runtime and administrative events, including status and error messages that can be used for troubleshooting. Audit information is logged synchronously to preserve transactional integrity.

Although the bulk of logging is executed asynchronously, decreasing the amount of information written to log files always provides the best possible performance.

Starting with version 8.2, PingFederate only records messages that are tagged with log level INFO, WARN, ERROR, and FATAL to the server log (and the provisioner log). Messages that are tagged DEBUG (or TRACE) are not recorded to optimize performance. Console logging is also disabled for the same reason.

For troubleshooting purpose, you may adjust the log level to DEBUG in the `log4j2.xml` file and optionally re-enable console logging.

 **Important:** When debug messages and console logging are no longer required, ensure they are turned off.

For more information, see [Enable debug messages and console logging](#) on page 513.

Operating system tuning

This section contains tuning recommendations for your operating system. The tuning recommendations provided here are particularly useful in preventing deployment issues in high capacity environments.

Linux tuning

Implement these recommendations in your Linux environment to prevent deployment issues, increase the performance and the capacity of the networking stack (particularly TCP) and the file descriptor usage, and thus enabling PingFederate® to handle a high volume of concurrent requests.

Network/TCP tuning

Add or modify the following entries in the `/etc/sysctl.conf` file:

```
##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

# Enables fast recycling of TIME_WAIT sockets.
# (Use with caution according to the kernel documentation!)
net.ipv4.tcp_tw_recycle=1

# Allow reuse of sockets in TIME_WAIT state for new connections
# only when it is safe from the network stack's perspective.
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:

```
pf_user soft nofile 10400
pf_user hard nofile 10400
```

where `pf_user` is the user account used to run the PingFederate java process (or `*` for all user accounts).

Windows tuning

Implement these recommendations in your Windows environment to prevent deployment issues, increase the performance and the capacity of the networking stack (specifically the TCP socket), and thus enabling PingFederate® to handle a high volume of concurrent requests.

Increase the number of available ephemeral ports

1. Start the Command Prompt application (`cmd.exe`).
2. Type `netsh int ipv4 show dynamicportrange tcp` to view the ephemeral ports.
3. Type `netsh int ipv4 set dynamicport tcp start=1025 num=64510` to increase the range of the ephemeral ports. (Administrative privileges required.)
4. Reboot the server.
5. Type `netsh int ipv4 show dynamicportrange tcp` to confirm the updated port range.

Reduce socket `TIME_WAIT` delay

1. Start the Registry Editor application (`regedit.exe`).
2. Navigate to `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters`.
3. Create a new **DWORD (32 bit) value** and set the name as `TcpTimedWaitDelay`.
4. Set a decimal value of 30.
5. Reboot the server.

Concurrency

Generally speaking, the more requests that are processed in parallel, the higher the number of requests that are 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

It is important to note that this topic serves as a guideline for optimizing the concurrency of your deployment. On a large system with multiple CPU (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 (reasonably expected) user load. This way, CPU resources are not under-utilized while still allowing room for occasional load spikes. The level of concurrency in PingFederate may 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 topic therefore serves as a guideline of where to start when tuning the server.

Acceptor thread pool

For optimal performance, particularly in larger deployments, PingFederate® uses a non-blocking I/O model to process requests since version 8.0.1. Administrators may fine-tune the queue size, as needed (see [step 5](#)).

For PingFederate 8.0, follow these steps to configure PingFederate 8.0 to use a non-blocking I/O model.

1. Stop PingFederate.

2. Edit the `<pf_install>/pingfederate/etc/jetty-runtime.xml` file.

Consider making a backup copy of this file.

3. Find the following section:

```
<Call id="httpsPrimaryConnector" name="addConnector">
  <Arg>
    <New
      class="com.pingidentity.appserver.jetty.server.connector.ServerConnector">
        <Arg name="server">
          <Ref refid="RuntimeServer"/></Arg>
          <Arg name="acceptors" type="int">2</Arg>
          <Arg name="selectors" type="int"><Property
            name="ssl.selectors" default="-1"/></Arg>
          ...
          <Set name="acceptQueueSize">512</Set>
        </New>
      </Arg>
    </Call>
```

4. Disable the acceptor threads and enable the non-blocking I/O model with the following updates (in **bold**):

```
<Call id="httpsPrimaryConnector" name="addConnector">
  <Arg>
    <New
      class="com.pingidentity.appserver.jetty.server.connector.ServerConnector">
        <Arg name="server">
          <Ref refid="RuntimeServer"/></Arg>
          <Arg name="acceptors" type="int"><Property
            name="ssl.selectors" default="0"/></Arg>
          <Arg name="selectors" type="int"><Property
            name="ssl.selectors" default="1"/></Arg>
          ...
          <Set name="acceptQueueSize">512</Set>
        </New>
      </Arg>
    </Call>
```

5. 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` element.
6. When finished, save your changes and restart PingFederate.
7. For a clustered PingFederate environment, repeat these steps on each engine node as needed.

Server thread pool

When tuning the server thread pool, it is best to size the system based on expected user load, setting 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. Testing has shown that PingFederate performs quite well when the server thread pool is sized between 25 and 50 server threads per available CPU core, assuming sufficient memory (see [Memory](#) on page 642). Note that this is guidance and should not necessarily be scaled directly 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

1. Stop PingFederate®.
2. Edit the `<pf_install>/pingfederate/etc/jetty-runtime.xml` file.

Consider making a backup copy of this file.

3. Find the following section:

```
<Get name="ThreadPool">
  <Set name="minThreads" type="int">10</Set>
  <Set name="maxThreads" type="int">200</Set>
  <Set name="detailedDump">false</Set>
</Get>
```

4. Modify it using the following guidelines:

- Set minThreads to (available CPU cores * 25)
- Set maxThreads to (available CPU cores * 50)

For example, if your PingFederate system has 1 CPU with 4 cores, the total available cores is 4. The configuration would be:

```
<Get name="ThreadPool">
  <Set name="minThreads" type="int">100</Set>
  <Set name="maxThreads" type="int">200</Set>
  <Set name="detailedDump">false</Set>
</Get>
```

5. When finished, save your changes and restart PingFederate.
6. For a clustered PingFederate environment, repeat these steps on each engine node as needed.

Connection pools to data stores

Database (JDBC) and LDAP data stores use connection pooling to improve the performance and efficiency of communicating with external systems. Connection pools improve efficiency by maintaining persistent connections to the database or LDAP server thus preventing the expense of creating the connection on demand. Connection pools also allow more control over the load placed on the back-end server.

Much like the server thread pool, for optimal performance, a number of connections is required to handle most (if not all) the requests in parallel. It may 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 database and LDAP resources. For optimal performance, connection pools should be sized large enough to handle between 50% and 100% of the number of concurrent requests the server is expected to encounter most often.

Of course, it is also very important to understand the capacity and limitation of the database or LDAP server and to size the connection pool accordingly. Sizing the connection pool beyond the capability of the back-end server could lead to PingFederate® flooding the data store without any performance improvement.

1. For JDBC data stores:

- a) On the **Server Configuration > Data Stores** screen, select the applicable JDBC data store.
- b) On the **Database Config** screen, click **Advanced**.
- c) On the **Advanced Database Options** screen:
 - 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 capability of the back-end database server.



Tip: The maxThreads value is defined in the `<pf_install>/pingfederate/etc/jetty-runtime.xml` file (see [Server thread pool](#) on page 640).

2. For LDAP data stores:

- a) On the **Server Configuration > Data Stores** screen, select the applicable LDAP data store.
- b) On the **LDAP Configuration** screen, click **Advanced**.
- c) On the **Advanced LDAP Options** screen:
 - Set the **Minimum Connections** value to 50% of the maxThreads value.

- Set the **Maximum Connections** value to between 75% and 100% of the maxThreads value, subject to the capability of the back-end database server.



Note: As of PingFederate 7.2, separate connection pools are created for bind requests and search requests. This dramatically improves performance of LDAP authentications, especially when using LDAPS, because bind requests no longer require unique connections. Each connection pool (bind and search) is configured individually per the values set for minimum and maximum connections

3. For a clustered PingFederate environment, replicate the changes to all engine nodes on the **Server Configuration > Cluster Management** screen.

Memory

After the CPU, memory is the most important resource in your deployment. The previous section describes how to configure PingFederate® to support more concurrent requests. With increased concurrency comes an increase in the memory required. Moreover, PingFederate is a Java application, and although this document is not to be considered a guide to garbage collection theory or ergonomics, consideration must be given to tuning that affects, or is affected by, garbage collection.

PingFederate will automatically set many of its own memory settings based on the available memory, not the total memory of the system.

JVM heap

The most important tuning that can be applied to the Java Virtual Machine (JVM) is the size of the heap memory. If the demands on PingFederate® 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 expensive processes, and thus detrimental to 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.

By default, the PingFederate launch script sets the size of the Java heap based on the amount of free memory detected at start-up time. In the case of running PingFederate as a Windows service, the heap size configuration is set based on the amount of free memory available when the service is installed.

- If less than 2 GB is free, the heap is set to a minimum of 256 MB, growing to a maximum of 1 GB. This means that when PingFederate is launched, 256 MB are available for newly created objects. If more than 256 MB are required, the JVM garbage collects and, if necessary, expands the heap until it reaches 1 GB. At this point, the JVM must invoke garbage collection in order to free up released memory to be re-used.
- If at least 2 GB of free memory is detected, the Java heap is set to 1.5 GB.
- If 2.5 GB or more of free memory is detected, the Java heap is set to 2 GB. Again, once the heap is full, the JVM will initiate garbage collection in order to free up memory to be re-used.

The JVM can grow the heap from the value of the minimum heap variable to the value of the maximum heap variable. However, growing the heap is not free. Requesting memory from the operating system is often an expensive exercise. In addition, the JVM must also reorganize the heap to account for the memory being added. If conserving memory in your deployment is important, then setting a lower value for the minimum heap than that of the maximum heap ensures that you are not reserving memory that you are not specifically using. If, however, you have enough memory such that a certain amount is easily earmarked for the PingFederate server, we recommend that you adjust the size of the heap by setting 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 being used exceeds the value of minimum heap.

Garbage collectors

The JVM comes with a wide variety of tuning options. Although the virtual machine should configure itself to run the best it can in most situations; it is sometimes advantageous to configure it for a specific application. This section explores a few additional options that can be applied to potentially improve the operation of PingFederate® in your

deployment. In all cases, it is recommended that you 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 are quite fast (see [Young generation bias](#) on page 643).

Concurrent mark sweep collector

The concurrent mark sweep (CMS) collector is suitable for 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 intend on using a JVM heap greater than 6 GB on larger machines with eight or more CPUs, the CMS collector may 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.


As needed, the CMS collector can be enabled by the `-XX:+UseConcMarkSweepGC` JVM option.

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).


If you intend to use the G1 collector, it is strongly advised that you remove any sizing options specific to the young generation. The reason is that 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, the G1 collector can be enabled by the `-XX:+UseG1GC` JVM option.

 **Tip:** For additional information about each garbage collector, please refer to [Java Platform, Standard Edition HotSpot Virtual Machine Garbage Collection Tuning Guide](#) from Oracle (docs.oracle.com/javase/8/docs/technotes/guides/vm/gctuning/).

Young generation bias

Without getting into too much detail on the generational memory management model in Java, the basic principal is that new objects are created in the *young generation* and are garbage collected when they are no longer used. If an object is created, and a reference is maintained, that object is eventually moved to the *old generation*.

 **Note:** If you intend to use the garbage first (G1) collector, let the JVM handles this aspect of the heap because specific settings can affect the performance of the G1 collector adversely.

Skip to [Configure JVM options](#) on page 644 for instructions on fine-tuning memory and garbage collection settings.

The processing model of PingFederate® is mostly geared towards short-lived transactions (single sign-on and token processing) and not long held, interactive, user sessions. As such, most of the objects that are created are relatively short-lived. It does not make sense to *promote* short-lived objects to the old generation because they are probably not going to be needed for long anyway. The problem is that when the young generation fills up, space must be made for new objects. So objects that are in the young generation and still in use must be moved to the old generation, which has a cost. Moreover, those objects will eventually 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 because the old generation is cleaned only during a full garbage collection (meaning that the JVM has almost reached the value of the maximum heap variable (-Xmx) and the entire heap must be cleaned). However, the young generation is garbage collected more frequently and with multiple threads by default (on systems with multiple cores), 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 total space of the heap. This means more frequently moving objects 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, the frequency of the more costly full collections is reduced. 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 (-Xmx) 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 are also applicable 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. Put another way, the size of the young generation is one fourth of the total heap size.

In a mostly short-lived object environment, it is recommended that 50-60% of the heap be given to the young generation; for example:

Young generation bias condition	JVM options
To fix a heap size of 2 GB with 50% the young generation bias using the <code>NewSize</code> argument	<code>-Xms2048m -Xmx-2048m -XX:NewSize=1024m - XX:MaxNewSize=1024m</code>
To fix a heap size of 2 GB with 50% the young generation bias using the <code>NewRatio</code> argument	<code>-Xms2048m -Xmx-2048m -XX:NewRatio=1</code>

Configure JVM options

PingFederate® attempts to optimize JVM heap settings based on available system resources, you can however override the default behaviors to apply specific heap sizing and other JVM options manually. Depending on the host operating system, you may tune the launch script or the Windows service configuration file.

PingFederate as a console application

When run as a console application on a Windows server or a Unix or Linux system, PingFederate® attempts to optimize JVM heap settings based on available system resources at start-up. The settings are defined in the launch scripts: the `run.bat` file (for Windows) and the `run.sh` file for Unix or Linux. Note that the `run.sh` file also controls the JVM settings when PingFederate is run as a non-interactive background process on a UNIX or Linux system. In essence, the launch scripts construct the JVM command line that launches the PingFederate process.

Any modifications or additions to the JVM command line should be applied to these files. JVM arguments are space-delimited. Arguments need not appear in any specific order; however, it is recommended that you edit specific entries in the script where noted.

It is also highly recommend that you make a backup copy of the launch script prior to making any changes so that you can revert to the default configuration as necessary. If you want to override the resource-based heap size or the garbage collector configuration determined at startup (or both), specify the values in the override section of the launch scripts.

1. Stop PingFederate.
2. Edit the launch script located in the <pf_install>/pingfederate/bin directory: the run.bat file (for Windows) or the run.sh file for Unix or Linux
3. To configure a specific heap size:
 - a) Find the following section:

run.bat (for Windows)

```
REM If you want to override
resource-based \
JVM configuration, you may do so
below.
...
REM For manual configuration, \
remove the preceding "REM" and \
apply values.
REM
REM set minimumHeap=-Xms256m
REM set maximumHeap=-Xmx1024m
REM set minimumPermSize=-
XX:PermSize=96m
REM set minimumNewSize=
REM set maximumNewSize=
REM set garbageCollector=
```

run.sh (for UNIX or Linux)

```
# If you want to override
resource-based \
JVM configuration, you may do so
below.
...
# For manual configuration, \
remove the preceding "#" and \
apply values between quotes.
#
# minimumHeap="-Xms256m"
# maximumHeap="-Xmx1024m"
# minimumPermSize="-
XX:PermSize=96m"
# minimumNewSize=""
# maximumNewSize=""
# garbageCollector=""
```

Note that line breaks (indicated by the backslash character, \) are inserted for readability.

- b) Remove the preceding REM or # for the minimumHeap and maximumHeap parameters; then override the default values.

The valid unit qualifiers are *k* for kilobytes, *m* for megabytes and *g* for gigabytes. That is, `maximumHeap=-Xmx1536m` and `maximumHeap=-Xmx1.5g` are equivalent.

For example, to fix the JVM heap size to 2 GB, set the minimum and maximum as follows:

run.bat (for Windows)

```
REM If you want to override
resource-based \
JVM configuration, you may do so
below.
...
REM For manual configuration, \
remove the preceding "REM" and \
apply values.
REM
set minimumHeap=-Xms2g
set maximumHeap=-Xmx2g
REM set minimumPermSize=-
XX:PermSize=96m
REM set minimumNewSize=
REM set maximumNewSize=
```

run.sh (for UNIX or Linux)

```
# If you want to override
resource-based \
JVM configuration, you may do so
below.
...
# For manual configuration, \
remove the preceding "#" and \
apply values between quotes.
#
minimumHeap="-Xms2g"
maximumHeap="-Xmx2g"
# minimumPermSize="-
XX:PermSize=96m"
# minimumNewSize=""
# maximumNewSize=""
```

run.bat (for Windows)

```
REM set garbageCollector=
```

run.sh (for UNIX or Linux)

```
# garbageCollector=""
```

4. To override the garbage collection option:

- Find the override section, as described in [step 3a](#).
- Remove the preceding REM or # for the garbageCollector parameter; then specify the option.

To override the number of garbage collection threads used by the parallel collector, use the `-XX:ParallelGCThreads=n` argument, where *n* is the desired number of garbage collector threads. For example, to set the number of garbage collector threads to 4, update as follows:

run.bat (for Windows)

```
set garbageCollector=-
XX:ParallelGCThreads=4
```

run.sh (for UNIX or Linux)

```
garbageCollector="-
XX:ParallelGCThreads=4"
```

To enable the concurrent mark sweep (CMS) collector, use the `-XX:+UseConcMarkSweepGC` argument; for example:

run.bat (for Windows)

```
set garbageCollector=-XX:
+UseConcMarkSweepGC
```

run.sh (for UNIX or Linux)

```
garbageCollector="-XX:
+UseConcMarkSweepGC"
```

To enable the garbage first (G1) collector, use the `-XX:+UseG1GC` argument; for example:

run.bat (for Windows)

```
set garbageCollector=-XX:+UseG1GC
```

run.sh (for UNIX or Linux)

```
garbageCollector="-XX:+UseG1GC"
```



Note: If you enable the G1 collector, it is recommended that you remove any sizing options specific to the young generation. Skip to [step 6](#).

5. To specify young generation-specific sizing options:

- Find the override section, as described in [step 3a](#).
- Remove the preceding REM or # for the minimumNewSize and maximumNewSize parameters; then enter the desired values.

Similar to the minimumHeap and maximumHeap parameters, the valid unit qualifiers are *k* for kilobytes, *m* for megabytes and *g* for gigabytes.

For example, to fix the young generation bias to 1 GB, set the minimum and maximum as follows:

run.bat (for Windows)

```
set minimumNewSize=1024m
set maximumNewSize=1024m
```

run.sh (for UNIX or Linux)

```
minimumNewSize="1024m"
maximumNewSize="1024m"
```

6. To remove young generation-specific sizing options:

- Find the override section, as described in [step 3a](#).
- Remove the preceding REM or # for the minimumNewSize and maximumNewSize parameters; then leave the value as-is (an empty value).

- To add additional arguments to the JVM command line, you must create a script variable to contain the text of the JVM argument and then append the new script variable to the script variable that makes up the final JVM command line.

For example, to enable the aggressive options flag update the launch script as follows:

run.bat (for Windows)

```
set aggOpts=-XX:+AggressiveOpts
set PF_JAVA_OPTS=%PF_JAVA_OPTS%
%aggOpts%
```

run.sh (for UNIX or Linux)

```
aggOpts="-XX:+AggressiveOpts"
JAVA_OPTS="$JAVA_OPTS $aggOpts"
```

8. When finished, save your changes and restart PingFederate.
9. For a clustered PingFederate environment, repeat these steps on each engine node as needed.

PingFederate as a Windows service

When PingFederate® is installed to run as a Windows service, PingFederate attempts to optimize JVM heap settings based on available system resources at the installation of the service. The settings are defined in the `PingFederateService.conf` file.

Any modifications or additions to the JVM command line should be applied to the `PingFederateService.conf` file. JVM arguments are space-delimited. Arguments need not appear in any specific order; however, it is recommended that you edit specific entries in the script where noted.

It is also highly recommend that you make a backup copy of the `PingFederateService.conf` file prior to making any changes so that you can revert to the default configuration as necessary.

- ❗ **Important:** Reinstalling the PingFederate Windows service creates a backup copy of the existing `PingFederateService.conf` file and creates a new `PingFederateService.conf` file with values based on the available resources at the time. In order words, you must re-customize any desired changes after the re-installation of the service.

1. Stop the PingFederate service.
2. Edit the `PingFederateService.conf` located in the `<pf_install>/pingfederate/sbin/wrapper` directory.
3. To configure a specific heap size:
 - a) Find the following section:

```
# Initial Java Heap Size (in MB)
wrapper.java.initmemory=256

# Maximum Java Heap Size (in MB)
wrapper.java.maxmemory=1024
```

- b) Modify the numerical values of the `wrapper.java.initmemory` and `wrapper.java.maxmemory` parameters in megabytes (MB) only.
4. To override the garbage collection option:
 - a) Find the following section:

```
# Java Additional Parameters
wrapper.java.additional.1=...
wrapper.java.additional.2=...
...
wrapper.java.additional.m=...
```

where *m* is the last additional parameter in the file.

- b) Add the garbage collection option as the `wrapper.java.additional.m+1` parameter.

- ❗ **Important:** When adding additional `wrapper.java.additional` parameters, ensure no parameters share the same numerical sequence number.

To override the number of garbage collection threads used by the parallel collector, use the `-XX:ParallelGCThreads=n` argument, where *n* is the desired number of garbage collector threads. For example, to set the number of garbage collector threads to 4, update as follows:

```
# Java Additional Parameters
wrapper.java.additional.1=...
wrapper.java.additional.2=...
...
wrapper.java.additional.16=-XX:ParallelGCThreads=4
```

To enable the concurrent mark sweep (CMS) collector, use the `-XX:+UseConcMarkSweepGC` argument; for example:

```
# Java Additional Parameters
wrapper.java.additional.1=...
wrapper.java.additional.2=...
...
wrapper.java.additional.16=-XX:+UseConcMarkSweepGC
```

To enable the garbage first (G1) collector, use the `-XX:+UseG1GC` argument; for example:

```
# Java Additional Parameters
wrapper.java.additional.1=...
wrapper.java.additional.2=...
...
wrapper.java.additional.16=-XX:+UseG1GC
```



Note: If you enable the G1 collector, it is recommended that you remove any sizing options specific to the young generation. Skip to [step 6](#).

5. To specify young generation-specific sizing options:
 - a) Find the override section, as described in [step 4a](#).
 - b) Add the young generation parameters.

Similar to the `wrapper.java.initmemory` and `wrapper.java.maxmemory` parameters, the values are in MB only.

For example, to fix the young generation bias to 1 GB, set the minimum and maximum as follows:

```
# Java Additional Parameters
wrapper.java.additional.1=...
wrapper.java.additional.2=...
...
wrapper.java.additional.17=-XX:NewSize=1024
wrapper.java.additional.18=-XX:MaxNewSize=1024
```

6. To remove young generation-specific sizing options:
 - a) Find the override section, as described in [step 4a](#).
 - b) Remove the following parameters (if found, and regardless whether or not values are assigned to the parameters):

```
wrapper.java.additional.17=-XX:NewSize=
wrapper.java.additional.18=-XX:MaxNewSize=
```

7. To add other JVM arguments, create additional `wrapper.java.additional.x` parameters and add the arguments as their values.

For example, to enable the aggressive options flag update the configuration file as follows:

```
# Java Additional Parameters
wrapper.java.additional.1=...
wrapper.java.additional.2=...
```

```
...
wrapper.java.additional.19=-XX:+AggressiveOpts
```

8. When finished, save your changes and restart PingFederate.
9. For a clustered PingFederate environment, repeat these steps on each engine node as needed.

Hardware security modules

For optimal security, PingFederate® can be configured 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, see [Supported hardware security modules](#) on page 51.)

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. It is recommended that you 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

For deployments with hundreds of connections or OAuth clients (or both), if noticeable delays are observed 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.

(For more information, see [Configure a backup schedule](#) on page 119 and [Configure automatic connection validation](#) on page 153, respectively.)

References

Memory management

[Java Platform, Standard Edition HotSpot Virtual Machine Garbage Collection Tuning Guide](https://docs.oracle.com/javase/8/docs/technotes/guides/vm/gctuning/) (docs.oracle.com/javase/8/docs/technotes/guides/vm/gctuning/)

Hotspot JVM arguments

[Java HotSpot VM Options](http://www.oracle.com/technetwork/java/javase/tech/vmoptions-jsp-140102.html) (www.oracle.com/technetwork/java/javase/tech/vmoptions-jsp-140102.html)

Index

A

- access control [108](#)
- account linking
 - configuring [347](#)
- account mapping [69](#)
- accounts, administrator [108](#)
- activation
 - for IdP connection [389](#)
 - for SP connection [327](#)
- adapters
 - configuring
 - IdP [275](#)
 - SP [331](#)
 - integration [67](#)
 - mapping URLs to SP [332](#)
- administrative console
 - certificate authentication to [112](#)
 - logging on using LDAP [112](#)
- administrators, adding [109](#)
- affiliations, SP [327](#)
- Apache Commons OGNL page [498](#)
- artifact
 - lifetime, setting for IdP connections [360](#)
 - lifetime, setting for SP connections [304](#)
- Artifact Resolution Service [37, 39](#)
- assertion [45, 47](#)
- Assertion Consumer Service [38](#)
- Assertion Consumer Service URLs [302](#)
- assertions
 - content [291](#)
 - lifetime [290](#)
 - mapping attributes to [296](#)
- Assertions [30](#)
- attribute authority (XASP), configuring [365](#)
- attribute contract
 - default for SP connection [298](#)
 - IdP connection [348](#)
 - to SP [292](#)
- attribute query
 - configuring for IdP connection [365](#)
 - configuring for SP connection [306](#)
 - mapping names [365](#)
- attributes
 - mapping
 - about [69](#)
 - for SP connection [297, 308, 409, 416](#)
 - masking in log files [74, 95](#)
 - overview [70](#)
- attributes:
 - sources for express provisioning [366](#)
- attribute sources
 - custom [165](#)
 - failover (for IdP) [296](#)
 - for IdP [296](#)
- audit logging [83](#)
- Auto-Connect
 - about [79](#)

- activation (IdP connection) [399](#)
- activation (SP connection) [330](#)
- allowed IdP domains [399](#)
- allowed SP domains [330](#)
- configuring (as an IdP) [329](#)
- configuring (as an SP) [398](#)
- metadata lifetime [157](#)
- metadata signing [156](#)

B

- back-channel settings
 - for IdP connections [385](#)
 - for SP connections [310](#)
- bindings
 - allowable
 - for IdP connection [359](#)
 - for SP connection [304](#)
 - configuring
 - for IdP connection [347](#)
 - for SP connection [290](#)
 - SOAP for IdP connection [385](#)
 - SOAP for SP connection [310](#)
- buttons, administrative console [29](#)

C

- certificate authority (CA) [178](#)
- certificates
 - authority, verifying [78](#)
 - considerations [83](#)
 - digital signing [184](#)
 - expiration notification [146](#)
 - exporting [184](#)
 - for SOAP authentication [83](#)
 - importing [184](#)
 - management [75](#)
 - revocation of [191](#)
 - selecting XML decryption (SP-to-IdP) [389](#)
 - selecting XML encryption (SP-to-IdP) [388](#)
 - signature verification [184](#)
 - SSL [146, 179](#)
 - SSL client [182](#)
 - trust models for [77](#)
 - verifying [78](#)
- checklist, for federation [82](#)
- common domains [175](#)
- Common Event Format [105](#)
- connections
 - importing [477](#)
 - migrating [120](#)
- console buttons [29](#)
- CSR, importing/exporting [184](#)

D

- data exchange
 - automated [85](#)
 - for IdP connection [85](#)
 - for SP connection [85](#)
- data store
 - compatibility [14](#)
 - description [72](#)
 - for attribute query profile [307](#)
 - introduction [158](#)
 - JDBC configuration [159](#)
 - LDAP configuration [161](#)
 - selecting
 - for express provisioning [155, 367](#)
 - for SaaS provisioning [320](#)
 - troubleshooting [514](#)
- data store:selecting
 - for express provisioning [155, 367](#)
- decryption keys
 - for IdP connections [389](#)
 - for SP connections [315](#)
- defederation [70](#)
- DeficientPartnerWorkaround [143, 586](#)
- deployment diagrams [13](#)
- digital signatures
 - about [75](#)
 - keys and certificates [184](#)
 - policy [77](#)
- digital signing
 - for IdP connection [387](#)
 - for SP connection [313](#)
- directory service
 - code example [481](#)
 - SOAP example [482](#)
- discovery, IdP [174](#)
- domains
 - allowed for IdP Auto-Connect [399](#)
 - allowed for SP Auto-Connect [330](#)

E

- email
 - addresses [146](#)
 - notifications [146](#)
 - server configuration [114](#)
- encryption, XML [305, 364](#)
- endpoints
 - protocol, SP [337](#)
- endpointsapplication endpoints
 - SP [336](#)
 - SP application [336](#)
- entity ID [152](#)
- event trigger, for provisioning [372](#)

F

- failover
 - for data stores (IdP) [296](#)
- federation
 - checklist [82](#)
 - choosing [151](#)

- defederation [70](#)
- ID [85](#)
- roles [151](#)
- URLs [152](#)
- for XASP [337](#)

H

- hardware requirements [14](#)
- healthcare.gov [143, 586](#)
- heartbeat, server
 - checking via HTTP [460](#)
- host names, virtual [115](#)

I

- identity mapping
 - about [69](#)
 - IdP configuration [291](#)
 - SP configuration [347](#)
- Identity Provider [30](#)
- IdP
 - adapters [275](#)
 - connections
 - activation [389](#)
- IdP Discovery
 - common domain service [175](#)
 - configuring [174](#)
 - local domain service [175](#)
 - selecting [151](#)
- importing connections [477](#)
- importing metadata
 - for IdP connection [342](#)
 - for SP connection [286](#)
- installation
 - federation server [19](#)
 - license key [91](#)
 - Linux service [25](#)
 - Server JRE [19](#)
 - Windows service [25](#)
- installing
 - java [19](#)
- integration kits [67](#)

J

- java [19](#)
- Java requirements [14](#)
- JDBC configuration [159](#)
- JRE [19](#)

K

- keys
 - for XML decryption (SP-to-IdP) [389](#)
 - for XML encryption (IdP-to-SP) [315](#)
- keys and certificates
 - digital signature [184](#)
 - SSL [182](#)

L

- LDAP
 - configuration [161](#)
 - SSL [161](#)
 - using SSL [161](#)
- LDAP, using for authentication [112](#)
- license key, installing [91](#)
- licensing
 - email address to notify [146](#)
- Linux
 - service [25](#)
- log files [95](#)
- logging
 - administrator audit [97](#)
 - files [95](#)
 - transaction [99](#)
 - using Common Event Format [105](#)

M

- mapping
 - adapters
 - URLs to SP [332](#)
 - attributes
 - configuring for SP connection [297](#), [308](#), [409](#), [416](#)
 - identity [69](#), [291](#), [347](#)
- masking attributes [74](#)
- masking attributes in logs [95](#)
- message
 - signing [75](#)
 - validation [76](#), [83](#)
- metadata
 - exporting [116](#)
 - importing
 - for IdP connection [342](#)
 - for SP connection [286](#)
 - providing for Auto-Connect [79](#)
 - signing [118](#)
 - signing for Auto-Connect [156](#)
 - use [84](#)
- migrating connections [120](#)
- monitoring
 - JMX [148](#)
 - provisioning [148](#)
 - server [147](#)

N

- name identifiers
 - SAML 2.0 [291](#)
 - WS-Federation [292](#)
- navigation
 - buttons [29](#)
- notificationslicensing
 - expiration notification [146](#)
 - runtime [146](#)

O

- OAuth
 - enabling [242](#)
- Object-Graph Navigation Language (OGNL), editing [501](#)
- OpenID Provider [30](#)
- operating systems [14](#)

P

- password notification [108](#)
- PF-12489 [143](#), [586](#)
- PF-16514 [143](#)
- PF-16860 [143](#)
- planning checklist [82](#)
- profiles
 - configuring
 - for SP connection [290](#)
 - for IdP connection [347](#)
- protocol endpoints
 - SP [337](#)
- protocols
 - troubleshooting [522](#)
- protocolsstandards
 - choosing [151](#)
- provisioning
 - enabling [347](#)
 - event trigger for [372](#)
- provisioning, as SP
 - choosing data stores for [367](#)
- provisioning, SaaS
 - choosing internal data store [155](#)
 - choosing source data stores for [320](#)
- pseudonym [45](#)
- pseudonyms
 - unique values for [276](#)

R

- Relying Party [30](#)
- requirements
 - data store [14](#)
 - hardware [14](#)
 - Java [14](#)
 - operating systems [14](#)
- Roles [30](#)
- runtime
 - configuration [93](#)
 - reporting [147](#)

S

- SAML
 - selecting [151](#)
- SDK [68](#)
- security
 - back channel [83](#)
 - considerations [75](#)
- Serve JRE installation [19](#)
- server
 - troubleshooting [514](#)

- server HTTP checking [460](#)
- server ID [83](#)
- server monitoring [147](#)
- server status verification [460](#)
- Service Provider [30](#)
- service URL
 - IdP, for WS-Fed [358](#)
 - SP, for WS-Fed [303](#)
- session integration considerations [83](#)
- signature verification [184](#)
- signing XML files [118](#)
- SLO service URLs [304](#)
- SNMP configuration [147](#)
- SOAP [37](#)
- SOAP responder URL [85](#)
- SP
 - adapters [331](#)
 - connections
 - activation [327](#)
 - default URL [336](#)
- SSL
 - about [78](#)
 - certificates [179](#)
 - keys and certificates [182](#)
 - using for LDAP [161](#)
- SSO [480](#)
- SSO directory service [480](#)
- SSPR [586](#)
- starting and stopping the server
 - Linux service [25](#)
- summary
 - for IdP connection [389](#)
 - for SP connection [327](#)
- synchronization
 - clock [83](#)
 - provisioning [155](#)
- system administration (single or multiple users) [146](#)
- system requirements [14](#)

T

- time synchronization [83](#)
- transaction logging [83, 99](#)
- transport security considerations [83](#)
- trigger, for provisioning [372](#)
- troubleshooting
 - data store [514](#)
 - protocols [522](#)
 - server [514](#)
- trust models [77](#)

U

- uninstalling [27](#)
- unique IDs [83](#)
- unique values, for pseudonym creation [276](#)
- user management
 - setting [146](#)
- users
 - account management [108](#)
 - adding [109](#)

V

- validating messages [83](#)
- verifying certificates [78](#)
- virtual host names [115](#)
- VirtualIdentityUtil [143, 586](#)
- virtual IDs [83](#)

W

- Web Services
 - connection management [477](#)
 - SSO Directory [480](#)
- Web templates [126](#)
- Windows service, installation [25](#)
- WS-Federation
 - setting realm ID [152](#)
 - SP service URL [303, 358](#)
- WS-Federationroles, choosing IdP/SP
 - selecting [151](#)
- WSS [48](#)
- WS-Trust STS
 - configuring request parameters [405](#)
 - enabling [400](#)
 - IdP connections [414](#)
 - mapping request parameters [407](#)
 - SP connections [405](#)

X

- X.509 certificates [75](#)
- XASP:requester mapping
 - requester mapping [337](#)
- XML encryption
 - about [78](#)
 - IdP-to-SP [305](#)
 - SP-to-IdP [364](#)
- XML files, signing [118](#)