IBM Security Directory Integrator
Version 7.2

Federated Directory Server Administration Guide

SC27-6211-00

IBM Security Directory Integrator
Version 7.2

Federated Directory Server Administration Guide

SC27-6211-00

Note
Before using this information and the product it supports, read the general information under "Notices" on page 73.

Edition notice

Note: This edition applies to version 7.2 of *IBM Security Directory Integrator* **licensed program (5724-K74) and to all subsequent releases and modifications until otherwise indicated in new editions.**

Contents

Figures

About this publication

IBM® Security Directory Integrator is an integrated development environment and runtime service for general-purpose, multi-format, multi-directional, real-time data movement, synchronization, and transformation.

IBM Security Directory Integrator Version 7.2 Federated Directory Integrator Administration Guide contains information about using Federated Directory Server console to design, implement, and administer data integration solutions.

It also contains information about using the System for Cross-Domain Identity Management (SCIM) protocol and interface for identity management.

Access to publications and terminology

Read the descriptions of the IBM Security Directory Integrator Version 7.2 library and the related publications that you can access online.

This section provides:

- A list of publications in the "IBM Security Directory Integrator library."
- Links to "Online publications" on page viii.
- A link to the "IBM Terminology website" on page ix.

IBM Security Directory Integrator library

The following documents are available in the IBM Security Directory Integrator library:

- *IBM Security Directory Integrator Version 7.2 Federated Directory Integrator Administration Guide*

  Contains information about using the Federated Directory Server console to design, implement, and administer data integration solutions. Also contains information about using the System for Cross-Domain Identity Management (SCIM) protocol and interface for identity management.
- *IBM Security Directory Integrator Version 7.2 Getting Started Guide*

  Contains a brief tutorial and introduction to IBM Security Directory Integrator. Includes examples to create interaction and hands-on learning of IBM Security Directory Integrator.
- *IBM Security Directory Integrator Version 7.2 Users Guide*

  Contains information about using IBM Security Directory Integrator. Contains instructions for designing solutions using the Security Directory Integrator designer tool (the Configuration Editor) or running the ready-made solutions from the command line. Also provides information about interfaces, concepts and AssemblyLine creation.
- *IBM Security Directory Integrator Version 7.2 Installation and Administrator Guide*

  Includes complete information about installing, migrating from a previous version, configuring the logging functionality, and the security model underlying the Remote Server API of IBM Security Directory Integrator. Contains information on how to deploy and manage solutions.
- *IBM Security Directory Integrator Version 7.2 Reference Guide*

Contains detailed information about the individual components of IBM Security Directory Integrator: Connectors, Function Components, Parsers, Objects and so forth – the building blocks of the AssemblyLine.

- *IBM Security Directory Integrator Version 7.2 Problem Determination Guide*

  Provides information about IBM Security Directory Integrator tools, resources, and techniques that can aid in the identification and resolution of problems.

- *IBM Security Directory Integrator Version 7.2 Message Guide*

  Provides a list of all informational, warning and error messages associated with the IBM Security Directory Integrator.

- *IBM Security Directory Integrator Version 7.2 Password Synchronization Plug-ins Guide*

  Includes complete information for installing and configuring each of the five IBM Password Synchronization Plug-ins: Windows Password Synchronizer, Sun Directory Server Password Synchronizer, IBM Security Directory Server Password Synchronizer, Domino® Password Synchronizer and Password Synchronizer for UNIX and Linux. Also provides configuration instructions for the LDAP Password Store and JMS Password Store.

- *IBM Security Directory Integrator Version 7.2 Release Notes®*

  Describes new features and late-breaking information about IBM Security Directory Integrator that did not get included in the documentation.

Online publications

IBM posts product publications when the product is released and when the publications are updated at the following locations:

IBM Security Directory Integrator Library
The product documentation site (http://pic.dhe.ibm.com/infocenter/tivihelp/v2r1/topic/com.ibm.IBMDI.doc_7.2/welcome.htm) displays the welcome page and navigation for this library.

IBM Security Systems Documentation Central
IBM Security Systems Documentation Central provides an alphabetical list of all IBM Security Systems product libraries and links to the online documentation for specific versions of each product.

IBM Publications Center
The IBM Publications Center site (http://www-05.ibm.com/e-business/linkweb/publications/servlet/pbi.wss) offers customized search functions to help you find all the IBM publications you need.

Related information

Information related to IBM Security Directory Integrator is available at the following locations:

- IBM Security Directory Integrator Version 7.2 uses the JNDI client from Oracle. For information about the JNDI client, see the *Java Naming and Directory Interface™ Specification* at http://download.oracle.com/javase/7/docs/technotes/guides/jndi/index.html .
- Information that might help to answer your questions related to IBM Security Directory Integrator can be found at https://www-947.ibm.com/support/entry/myportal/overview/software/security_systems/tivoli_directory_integrator.

IBM Terminology website

The IBM Terminology website consolidates terminology for product libraries in one location. You can access the Terminology website at http://www.ibm.com/software/globalization/terminology.

Accessibility

Accessibility features help users with a physical disability, such as restricted mobility or limited vision, to use software products successfully. With this product, you can use assistive technologies to hear and navigate the interface. You can also use the keyboard instead of the mouse to operate all features of the graphical user interface.

For additional information, see the Accessibility Appendix in *IBM Security Directory Integrator Version 7.2 Users Guide*.

Technical training

For technical training information, see the following IBM Education website at http://www.ibm.com/software/tivoli/education.

Support information

IBM Support provides assistance with code-related problems and routine, short duration installation or usage questions. You can directly access the IBM Software Support site at http://www.ibm.com/software/support/probsub.html.

IBM Security Directory Integrator Version 7.2 Problem Determination Guide provides details about:

- What information to collect before contacting IBM Support.
- The various methods for contacting IBM Support.
- How to use IBM Support Assistant.
- Instructions and problem-determination resources to isolate and fix the problem yourself.

Statement of Good Security Practices

IT system security involves protecting systems and information through prevention, detection and response to improper access from within and outside your enterprise. Improper access can result in information being altered, destroyed, misappropriated or misused or can result in damage to or misuse of your systems, including for use in attacks on others. No IT system or product should be considered completely secure and no single product, service or security measure can be completely effective in preventing improper use or access. IBM systems, products and services are designed to be part of a comprehensive security approach, which will necessarily involve additional operational procedures, and may require other systems, products or services to be most effective. IBM DOES NOT WARRANT THAT ANY SYSTEMS, PRODUCTS OR SERVICES ARE IMMUNE FROM, OR WILL MAKE YOUR ENTERPRISE IMMUNE FROM, THE MALICIOUS OR ILLEGAL CONDUCT OF ANY PARTY.

Chapter 1. Federated Directory Server

Federated Directory Server enables a collection of directories and other sources of data to be combined and treated as a single hierarchical directory. The Federated Directory Server console is a ready-to-use application that implements this directory integration.

Overview

IBM Security Directory Integrator provides extensive capabilities for developing complex data integration solutions. The Federated Directory Server console builds upon these capabilities and provides a quick and easy solution for connecting and synchronizing data from various sources.

IBM Security Directory Server is the core centralized repository for Federated Directory Server. The Federated Directory Server console provides synchronization services from one or more source systems, such as Active Directory or Sun Directory to the target, which is IBM Security Directory Server.

The Federated Directory Server console has the following advantages:

- Requires less implementation time and effort than custom-built solutions because it is a ready-to-use, quality application.
- Does not require in-depth knowledge of IBM Security Directory Integrator; hence it is easy to deploy and use.
- Enables integration across various data sources such as directories, databases, legacy data, and flat files, without affecting existing systems.
- Facilitates rapid deployment of identity and access management applications through a single point of access.
- Offers high speed, scalable performance, and superior security.

Features

Federated Directory Server has several features that you help you quickly and easily implement directory integration solutions.

- Directory integration is possible without requiring changes to existing legacy data.
- It pulls data automatically into IBM Security Directory Server.
- All relationships can contain advanced mapping and data transformation.
- Both users and groups can be integrated.
- Directory hierarchies can be maintained or flattened.
- Groups, including dynamic groups, can be created in a Federated Directory Server implementation that spans sources.
- Enriched data about people can be created from linked and augmented data from multiple sources.
- Federated Directory Server can be configured so that the user authentication goes directly to the existing backend local systems. Password replication, which is considered a major cost, is not required.
- Search is enabled across all content that is in the existing directory and data infrastructure.

- Users can log in to the enterprise directory by using a unique attribute like email or employee ID.
- Legacy data and the custom mapping of attributes can be managed through an interface that is easy to use.
- Write-back can be enabled to update the original sources.

Business scenarios

Federated Directory Server is a hybrid approach that addresses the security and collaboration requirements of directory services in various business scenarios.

The following examples are some of the business needs that the features of Federated Directory Server can fulfil:

- You want to enable a central authentication service. However, you might want to leave passwords in place in the original source directory.
- You are required to manage groups across multiple directories to support services like enterprise messaging and access control.
- You want to augment your identity information so that the central LDAP directory can support the specific needs of applications and services.

While IBM Security Directory Server is the centralized core back-end directory server, Federated Directory Server treats it more like a cache of information. Unless you want to do so, you do not have to use IBM Security Directory Server to manage the data. You can choose the level of service that you require.

The specific needs of customers can be categorized into the following scenarios that are illustrated in the diagram.

Migrate directories or co-exist

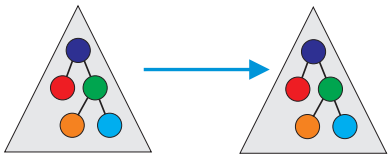

You can define the schemas and the amount of information that must be migrated. For example, you can provide more scalability and flexibility to data sources by migrating to Federated Directory Server without having to expand the schemas in the original data source.

Merge several data sources or directories

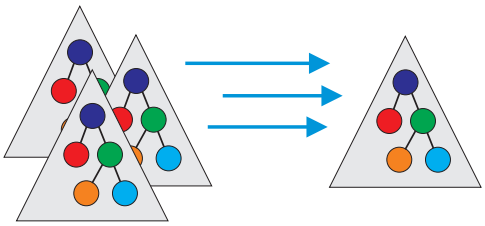

When you migrate or merge data from different data sources, the relationships can contain advanced mapping and data transformation. For example, you can integrate users and groups, maintain or flatten directory hierarchies, and create dynamic groups in Federated Directory Server that span data sources.

Enrich or augment with data from other sources

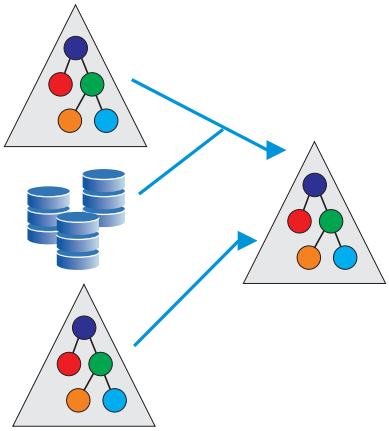

You can selectively add more data with a particular condition from another data source by setting up a join with the endpoint.

Selectively write back changes to the original source

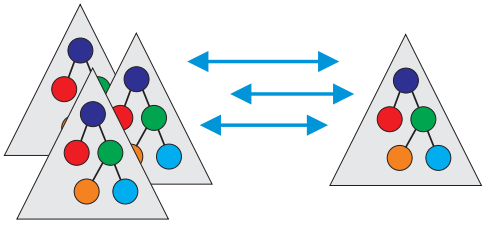

If information is modified in the target directory server, it can be written back to the endpoint. However, the write-back selective because some customers might want a barrier to preserve the original data in the endpoint.

Federate authentication back to the original source

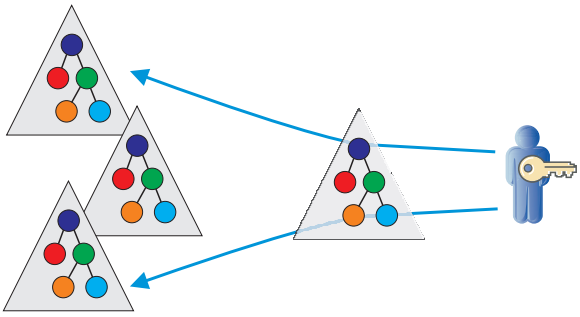

Federated Directory Server can send the authentication request back to the endpoint where the credentials are stored so that the authentication process happens at the endpoint. The credentials are not required to be stored in the Federated Directory Server unless you choose to do so.

For example, you can combine the various capabilities of Federated Directory Server to create a custom solution that is specific to your requirements. Assume that you have an Active Directory that you want to use for single sign-on. You want to provide more scalability to it for more uses like social networking, but do not want to expand the schemas. You can migrate the data selectively, for example,

only the email addresses of the users. Federated Directory Server also pulls the distinguished name (DN) from the source directory. You can then use the pass-through authentication capability of Federated Directory Server and retain the password credentials in the source directory itself without pulling it into the target directory. The user can log in to IBM Security Directory Server by using a unique attribute, which is the email address in this case. IBM Security Directory Server does a bind with the DN back to the Active Directory from where the user came. If a successful response is returned, then the user is authenticated.

Functional overview

Understand the key concepts, components, and architecture of Federated Directory Server.

The following diagram illustrates the various components of Federated Directory Server, which are described next.

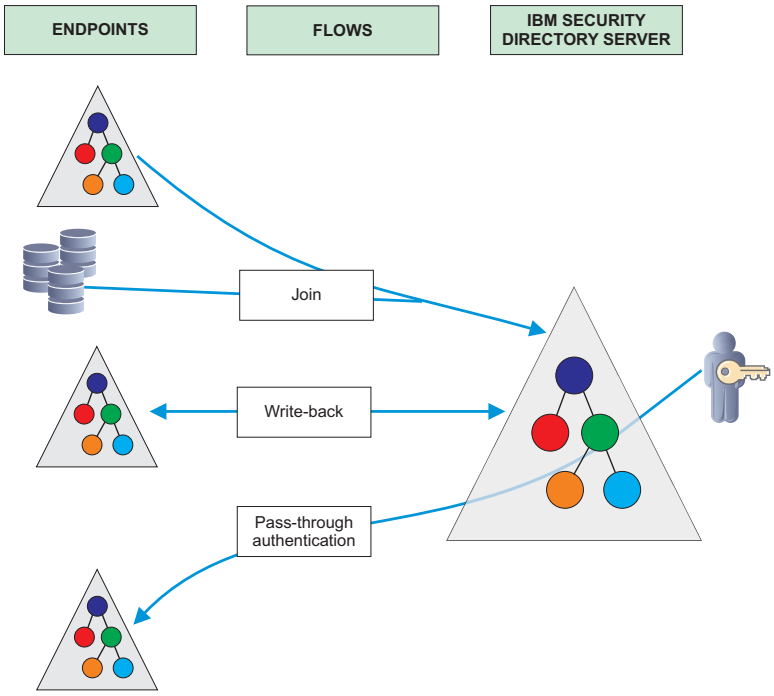

Figure 1. Federated Directory Server components

Directory Server
The IBM Security Directory Server, which is the target for all flows in the project.

Endpoint
A configured source system that can provide data in a flow. The endpoint types that are currently available are Active Directory, Custom AssemblyLine, File, JDBC, LDAP, IBM Security Directory Server and Sun Directory.

Flow A configuration that defines the relationship between the endpoints and the target IBM Security Directory Server. You must create flows only after you configure the target Directory Server connection settings and add one or more endpoints.

Attribute maps
A map that is used to convert the attribute from the source schema to the corresponding attribute in the target schema. In Federated Directory Server you can apply one of the ready-to-use attribute maps or a customized attribute map to a flow operation.

Join A configured source system that provides data that augments and enriches the data from the endpoint. If you configure a flow to specify a join with the endpoint, the entries are processed in the following manner:

1. An entry comes in from the endpoint.
2. The flow looks it up on the join data source.
3. The entry is merged with the data from the endpoint.
4. The merged data is added to the target directory server.

Pass-through authentication
A feature of IBM Security Directory Server where a user can be authenticated by delegating authentication to a different LDAP server. A flow contains a section for pass-through authentication. When you enable pass-through authentication for a flow, it configures IBM Security Directory Server to use the credentials that are stored in the endpoint for authenticating users that originate from that flow.

Roadmap for getting started

Use the roadmap to understand the key tasks for setting up your Federated Directory Server configuration and to run synchronization operations.

Table 1. Roadmap for getting started with Federated Directory Server

Key steps	Optional or advanced tasks
Understand the key concepts, components, and architecture.	
Access the Federated Directory Server console.	Configure security settings for accessing the console.
Connect to the target directory server.	Define custom attribute mapping between source endpoint and target directory server. Specify the log settings for the directory server.
Add endpoints and configure them for one or more of the following data sources: • LDAP • Active Directory • IBM Security Directory Server • Sun Directory • JDBC • File and file parsers • custom AssemblyLine (AL)	
Create a flow to define the relationship between the endpoints and the target directory.	

Table 1. Roadmap for getting started with Federated Directory Server (continued)

Key steps	Optional or advanced tasks
Define flow settings.	Extend the custom attribute map for a specific flow. Configure a join to augment the data from another data source selectively. Configure pass-through authentication to delegate authentication back to the endpoint. Enable write-back to propagate changes that are made in the target directory server back to the endpoint.
Verify the flow configuration by running a simulated synchronization operation.	
Run the initial synchronization to migrate data to the target directory.	
Schedule periodic incremental synchronizations.	Manually run a synchronization operation.
Use logs and reports to troubleshoot the flow configuration and synchronization operations.	Check the known issues and limitations to resolve specific issues.

Accessing the Federated Directory Server console

You can access the web-based Federated Directory Server console application in your browser.

Before you begin

Install IBM Security Directory Integrator Version 7.2.

About this task

The Federated Directory Server console is installed when you install IBM Security Directory Integrator Version 7.2. The Federated Directory Server console artifacts and configuration files are copied to a folder named `LDAPSync` in the IBM Security Directory Integrator solution directory. The installation is completed when you access the Federated Directory Server console for the first time.

Procedure

1. Start IBM Security Directory Integrator either from your system's launch interface or from the command line with the **ibmditk** or **ibmdisrv** command.
2. Open the Federated Directory Server console.
   - If you are accessing Federated Directory Server from your local system, click **Start** > **Programs** > **IBM Security Directory Integrator 7.2** > **Federated Directory Server Console**.
   - If you are accessing Federated Directory Server from a remote system, open the following link in your browser:

     `https://`*`hostname`*`:1098/fds`
3. If the security settings to access the Federated Directory Server console indicate that authentication is required, a login screen is displayed.

- If you are accessing from the `localhost`, specify the user name as `admin` and the password as `admin`, which are the defaults, and then click **Login**.
- If you are accessing from a remote system, you must specify the appropriate authentication credentials according to the security settings specified by the administrator. For more information, see "Security settings."

When the Federated Directory Server console is opened for the first time, the default configuration file named `LDAPSync.xml` is copied to the `configs` folder of the IBM Security Directory Integrator solution directory. A default project is created automatically, which you can use to configure Federated Directory Server.

What to do next

The default project does not have any endpoints or flows. To configure the project, you must complete the following steps:

1. Connect to a target directory server.
2. Configure one or more endpoints.
3. Define flow settings.

Note: As you configure the various features of Federated Directory Server in the console, by default, the changes are saved automatically. You can modify the autosave and refresh settings for the console:

1. On the Federated Directory Server console menu bar, click **Options**.
2. If you want to manually save the configuration changes that you make in the console, clear the **Enable auto-save** check box.
3. If you do not want to automatically reload the configuration changes, clear the **Automatically update FDS when configuration is saved** check box.
4. To create a snapshot of the current configuration before you make further changes, specify a **Snapshot description** and then click **Create snapshot**.
5. You can later roll back the changes to the level when you created a snapshot. Select the snapshot from the **Load snapshot** and then click **Load**.

Security settings

Access to the Federated Directory Server console is controlled by a set of properties that specify the security settings.

You must specify the security settings in the `solution.properties` file in the IBM Security Directory Integrator solution directory. These properties control the access to all of IBM Security Directory Integrator web applications, such as the Dashboard, REST API, and Federated Directory Server console.

Local and remote users are distinguished by the client IP address in the incoming access request:

- If the IP address belongs to one of the network cards on the system where IBM Security Directory Integrator is running, it is considered a `localhost` user.
- All other IP addresses are considered as remote users.

Access permission for `localhost` users is built in with the following credentials:

User name: `admin`

Password: `admin`

To specify access control and permissions, you can set or modify the following authentication properties:

dashboard.auth=true
Indicates whether users are required to authenticate.

Valid values are `true` if users are required to authenticate or `false` if no authentication is required.

dashboard.auth.localhost
Indicates the type of authentication that connections from the `localhost` must use.

Valid values are:

- `properties` specifies that property-based authentication must be used.
- `none` specifies that authentication is not required.
- `deny` specifies that all connections from `localhost` are denied.
- `ldap` specifies that authentication is done by logging in to an LDAP server and optionally validating group membership.

dashboard.auth.remote
Indicates the type of authentication that remote connections must use.

Valid values are:

- `properties` specifies that property-based authentication must be used.
- `none` specifies that authentication is not required.
- `deny` specifies that all remote connections are denied access, that is, all connections that are not from the `localhost` are denied access.
- `ldap` specifies that authentication is done by logging in to an LDAP server and optionally validating group membership.

{protect}-dashboard.auth.user.admin=admin
Specifies the user as `admin` with password `admin`.

dashboard.auth.ldap.url
Specifies the LDAP server address to use for authenticating the user. This property is used only if you specified `ldap` as the authentication mechanism.

Enter the LDAP host name, port number, and optionally a search base in the following format:

`ldap://`*`host`*`:`*`port`* `[/`*`search-base`*`]`

For example:

```
ldap://localhost:10389/ou=system
```

If the user provides an email address in the user name input field, IBM Security Directory Integrator first searches for a unique entry in the LDAP server from which it extracts the distinguished name (DN). Otherwise, it is expected that the value that is provided is acceptable to the LDAP server. After IBM Security Directory Integrator obtains a DN for the user name and the password from the user, it does an LDAP basic authentication with the DN and password.

dashboard.auth.ldap.url.group
Specifies the LDAP server address to use for verifying group membership of the user after authentication. This property is used only if you specified `ldap` as the authentication mechanism.

Enter the LDAP host name, port number, and optionally a search base in the following format:

`ldap://host:port [/search-base]`

For example:

```
ldap://localhost:389/cn=group1,ou=groups,ou=system
```

After the user is authenticated through LDAP, you can use this property to apply an additional group membership test before it allows access to the user.

You can also configure these properties in the IBM Security Directory Integrator Dashboard graphical user interface. In the Dashboard window, click **Actions** > **Show Server Details** > **Security and Connection**. For more information, see the IBM Security Directory Integrator documentation and search for *configuring Dashboard security settings*.

Connecting to IBM Security Directory Server

IBM Security Directory Server is the core centralized repository for Federated Directory Server. To use its synchronization services from one or more source systems to the target directory server, you must define the connection parameters for the target IBM Security Directory Server in the Federated Directory Server console.

Before you begin

Ensure that you select the project in which you want to specify the configuration settings. See "Accessing the Federated Directory Server console" on page 6.

Procedure

1. In the Federated Directory Server console navigation pane, under **Directory Server**, click **Connection Settings**.
2. On the **Connection Settings** page, in the **LDAP URL** field, enter the details of the target IBM Security Directory Server.

   The LDAP URL format is `ldap://hostname:port` or `ldap://server_IP_address:port`.

   For example: `ldap://localhost:389`.
3. In the **User Login** and **Password** fields, enter the distinguished name and credentials for authentication to the server.
4. In the **Default Target Container** specify the container in the target IBM Security Directory Server that is used to store the synchronized data.
5. In the next field, You can also specify a list of attributes that must be treated as binary, for example, **jpegPhoto**. The format is one attribute name on each line.
6. To verify the connection settings, click **Test Connection**.

What to do next

Specify the log settings for the target IBM Security Directory Server.

Specifying the log settings

After you configure the connection settings for IBM Security Directory Server, you can specify the path for the log file and log settings.

Procedure

1. In the Federated Directory Server console navigation pane, under **Directory Server**, click **Log Settings**.
2. On the **Log Settings** page, in the **Log Directory** field, you can specify the path for the log files. The default path is `LDAPSync/logs`.

   Note:
   - You can specify a path that is relative to the solution directory or current working directory of IBM Security Directory Integrator.
   - You can use forward slashes so that it is applicable to both Windows and UNIX systems.
3. In the **Log File History** field, specify the number of previous log files that must be retained. The default value is `20`.

What to do next

Configure one or more data resources as endpoints. See the following topics for the steps to configure the different types of endpoints.

Customizing attribute maps

When data is federated from multiple sources, the attributes must be mapped correctly when they are synchronized with the single target directory. You can specify how to convert attributes from the source endpoint schema to the target schema by defining custom maps for attributes.

About this task

The attribute mapping for standard schema such as Active Directory and Sun Directory is built in. Additionally, some ready-to-use custom maps are provided in Federated Directory Server. However, you might require to modify or extend these attribute maps or create new custom maps in some scenarios. For example, you might require custom maps if you use databases or files as your endpoint.

Procedure

1. In the Federated Directory Server console navigation pane, under **Directory Server**, click **Attribute Maps**.
2. On the Attribute Maps page, various attribute maps are provided for person, group, and container objects. Expand the type of attribute map that you want to customize.
3. If you want to create a new attribute mapping, click **Add Attribute**.
4. In the Add Attribute window, select the attribute from the list of attributes in the target directory server. A new row is displayed with the selected attribute name under the **Directory Server Attribute** column.
5. Under **Endpoint Attribute / Assignment**, specify the attribute name in the source endpoint that must map to the target attribute.
6. Double-click the endpoint attribute name to specify more settings for the attribute mapping.
   a. Select **Enabled** to use this attribute mapping for the endpoint.
   b. Click **Simple Assignment** or **Scripted Assignment** to specify the type of mapping. If you select **Scripted Assignment**, you can define the assignment by writing JavaScript code or by calling a function in the *Solution Directory*`\LDAPSync\customScript.js` file. For more information,

see the IBM Security Directory Integrator documentation at http://pic.dhe.ibm.com/infocenter/tivihelp/v2r1/topic/com.ibm.IBMDI.doc_7.2/welcome.html and search for *Scripting in IBM Security Directory Integrator*.

c. Specify whether you want this mapping to be used for all operations, or only when either modifying an entry or creating an entry.

7. To delete the mapping for a specific attribute, click the check box on that row. Then, click **Remove Attribute** and click **OK** when the confirmation message appears.
8. After you complete adding the attributes for a map, click **Save**. Unless you save each map that you edited, the changes are lost.
9. You can also duplicate the entire map and then extend it with your custom attribute mapping. Click **Duplicate Map**, enter a name for the new map file, and then click **OK**, when the confirmation message appears. A new attribute map with all the attribute mapping entries of the source map is created.
10. To delete an attribute map and all its entries, click **Delete Map** and then click **OK** when the confirmation message appears.

Results

You can later select this custom attribute map for use during flow operations when you define the flow specifications.

All attribute maps are stored in the *Solution Directory*\LDAPSync directory.

Configuring endpoints

You must specify endpoints for synchronization with the target IBM Security Directory Server. You can configure multiple LDAP directories, databases, files, or even subtrees as endpoints in the Federated Directory Server console.

Before you begin

Ensure that you specify the connections settings for the target IBM Security Directory Server. See "Connecting to IBM Security Directory Server" on page 9.

Procedure

1. To specify a new endpoint, in the Endpoints section of the navigation pane, click **Add**. The **Add Endpoint** window is displayed.
2. In the **Name** field, enter a name to identify the endpoint.
3. From the **Select endpoint type** list, select the appropriate type of endpoint. The following types of endpoints are available:
   - Active Directory
   - Custom AssemblyLine
   - File
   - JDBC
   - LDAP
   - Sun Directory
   - Security Directory Server

Note: After you create a configuration page for a specific type of endpoint, you cannot change it later. You must delete and create an endpoint again for the type of endpoint that you want to configure.

Results

The configuration page with endpoint parameters is displayed, which differs for each endpoint type.

In the navigation pane, a status icon is displayed next to each endpoint. You can click **Refresh** to see the latest status.

- A green dot is displayed soon after you create an endpoint and remains until you click **Test Connection** in the endpoint.
- After you test that the connection is successful, the green dot is replaced by a green tick mark.
- If the connection fails, a red cross mark is displayed.

What to do next

Configure the parameters for the endpoint. See the following topics for the different endpoint types.

If you want to delete an endpoint that you created and configured, follow these steps:

1. Under the **Endpoints** section of the navigation pane, right-click the name of the endpoint that you want to delete and then click **Delete**.
2. Click **OK** when the confirmation message appears.

Note: Flows that are based on an endpoint are also automatically deleted when you delete the endpoint.

Configuring an Active Directory endpoint

To configure an Active Directory as an endpoint, you must specify the LDAP URL, login name with credentials, the search base, and root suffix.

Before you begin

Ensure that you create an endpoint and specify the type as **Active Directory**. See "Configuring endpoints" on page 11.

Procedure

1. On the Active Directory endpoint configuration page, in the **LDAP URL** field, enter the LDAP URL of the Active Directory service you want to access.

   The LDAP URL format is `ldap://`*`hostname`*`:`*`port`* or `ldap://`*`server_IP_address`*`:`*`port`*.

   For example: `ldap://localhost:389`

   Note: The default LDAP port number is 389. If you are using SSL, the default LDAP port number is 636. For more information about setting up SSL for Active Directory connections, see the IBM Security Directory Integrator documentation at http://pic.dhe.ibm.com/infocenter/tivihelp/v2r1/topic/com.ibm.IBMDI.doc_7.2/welcome.html and search for *Microsoft Active Directory SSL configuration*.

2. In the **User Login** and **Password** fields, enter the distinguished name and credentials for authentication to the service.

   For example: `cn=administrator,cn=users,dc=your_domain,dc=com`

3. In the **Include entries from the following container** field, enter the search base of the source directory under which entries are read for synchronization. Alternately, you can click **Contexts** and select from the **LDAP Search Base** list and then click **OK**.

   For example: `dc=your_domain,dc=com`

   Note: For Active Directory, this value must be set to the root suffix of the domain controller; otherwise, delete modifications are not detected.

4. To verify the Active Directory connection settings, click **Test Connection**. If the connection is successful, the attributes in the endpoint are displayed in a separate pane. You can use the **Filter** field to search the attributes.

5. Optional: You can also configure the following advanced parameters. Expand the **Advanced** section to view these parameters.

   Page Size
   Specify the number of entries per page that must be returned by the request. The default value is 500.

   Seconds Before Timeout
   Specify the maximum number of seconds to wait for the next changed Active Directory object. The default value is 0.

   Seconds Between Polling
   Specifies the number of seconds to sleep between successive polls. The default value is 60.

   Change State Key
   Specifies the name of the key or parameter that stores the change detection iterator state. The state key is used between runs to remember the last changed that was processed. If synchronization was stopped for any reason, when it is restarted, it can pick up from where it stopped.

   The value of this key must be unique for each endpoint. If you do not set this parameter, a value is computed automatically to ensure uniqueness.

   Binary Attributes
   Specify a list of attributes that must be interpreted as binary values instead of strings. When you enter the attribute names in this field, enter one attribute per line and do not use any separators.

What to do next

After you configure the endpoint, you can create a flow to define the relationship between the endpoint and the target directory server.

Configuring a custom AssemblyLine endpoint

You can specify an AssemblyLine that you previously created in the Configuration Editor to be an endpoint in Federated Directory Server.

Before you begin

- Create an AssemblyLine by using the Configuration Editor. Ensure that the configuration XML file for the AssemblyLine project is copied to the `configs`

folder of the Solution Directory. For more information, see the IBM Security Directory Integrator documentation and search for *Configuration Editor*.

- In the Federated Directory Server console, ensure that you create an endpoint and specify the type as **Custom AssemblyLine**. See "Configuring endpoints" on page 11.

Procedure

1. On the Custom AssemblyLine endpoint configuration page, in the **SDI Project configuration name** field, enter the name of the IBM Security Directory Integrator project that contains the AssemblyLine.
2. In the following fields, specify the AssemblyLines that must be used for processing entries:
   - **AL to read Person entries**
   - **AL to read Group entries**
   - **AL to read Container entries**

   You must enter the following details of the AssemblyLine project that you created in Configuration Editor:
   - Name of the IBM Security Directory Integrator project that contains the AssemblyLine
   - Name of the AssemblyLine

   Use the following format:

   Project Name:/AssemblyLines/*AssemblyLine Name*

   For example, if your project is named `OS400` and it contains an AssemblyLine named `ReadUsers`, then you would enter:

   ```
   OS400:/AssemblyLines/ReadUsers
   ```
3. To verify the custom AssemblyLine endpoint connection settings, click **Test Connection**. If the connection is successful, the attributes in the endpoint are displayed in a separate pane. You can use the **Filter** field to search the attributes.

What to do next

After you configure the endpoint, you can create a flow to define the relationship between the endpoint and the target directory server.

Configuring a file endpoint

To configure a file as an endpoint, you must specify the file path, type of entry, and the file parser.

Before you begin

Ensure that you create an endpoint and specify the type as **File**. See "Configuring endpoints" on page 11.

Procedure

1. On the File endpoint configuration page, in the **File Path** field, enter the path of the file that you want to access.
2. From the **Type of Entry** list, select `person`, `group`, or `container`.
3. To verify the file connection settings, click **Test Connection**. If the connection is successful, the attributes in the endpoint are displayed in a separate pane. You can use the **Filter** field to search the attributes.

4. Optional: You can also configure the following advanced parameters. Expand the **Advanced** section to view these parameters.

   Timeout (in seconds)
   Specify a positive number to indicate the number of seconds to wait between operations before timeout occurs.

   Specify 0 (zero) to wait forever.

   If you select the **Lock file** option, the **Timeout** value instead specifies how long to wait to acquire the lock.

   Lock file
   Select this option to indicate that an exclusive lock is acquired for writing to the file. This lock prevents the file from being opened for writing by another instance of Federated Directory Server or any other program until the lock is released.
5. From the **Parser** list, select the name of the parser that you require to access the file. For more information about configuring each parser, see "File parsers" on page 34.

What to do next

After you configure the endpoint, you can create a flow to define the relationship between the endpoint and the target directory server.

Configuring a JDBC endpoint

To configure a JDBC connection as an endpoint, you must specify the JDBC URL, username and password, schema, table name, and type of entry.

Before you begin

Ensure that you create an endpoint and specify the type as **JDBC**. See "Configuring endpoints" on page 11.

Procedure

1. On the JDBC endpoint configuration page, in the **JDBC URL** field, enter the JDBC connection URL for the database that you want to access.

   The following examples are some typical URLs for various JDBC providers:

   IBM DB2®
   `jdbc:db2://hostname:port/dbname`

   Informix®
   `jdbc:informix-sqli://hostname:port/dbname:informixserver=Informix Server Name`

   Oracle `jdbc:oracle:thin:@hostname:1521:SID`

   Microsoft SQL Server
   `jdbc:sqlserver://hostname:1433;databasename=dbname;`

   Sybase
   `jdbc:sybase:Tds:hostname:port/`

   Derby `jdbc:derby://hostname:port/server_database_path;options`

   IBM solidDB®
   `jdbc:solid://hostname:port`
2. In the **JDBC Driver** field, enter the JDBC driver implementation class name.

The following examples are some typical driver implementation class names for various JDBC providers:

IBM DB2
`com.ibm.db2.jcc.DB2Driver`

Informix
`com.informix.jdbc.IfxDriver`

Oracle `oracle.jdbc.driver.OracleDriver`

Microsoft SQL Server
`com.microsoft.sqlserver.jdbc.SQLServerDriver`

Sybase
`com.sybase.jdbc3.jdbc.SybDriver`

Derby `org.apache.derby.jdbc.ClientDriver`

IBM solidDB
`solid.jdbc.SolidDriver`

For more information about JDBC drivers, see the IBM Security Directory Integrator documentation at http://pic.dhe.ibm.com/infocenter/tivihelp/v2r1/topic/com.ibm.IBMDI.doc_7.2/welcome.html and search for *Understanding JDBC Drivers*.

3. In the **Username** and **Password** fields, enter the login name and credentials to access the specified database.
4. In the **Schema** field, enter the schema from the table of the database that you want to use.
5. In the **Table name** field, enter the table or view for the operations.
6. From the **Type of Entry** list, select `person`, `group`, or `container`.
7. To verify the JDBC connection settings, click **Test Connection**. If the connection is successful, the attributes in the endpoint are displayed in a separate pane. You can use the **Filter** field to search the attributes.
8. Optional: You can also specify a custom `SELECT` statement to specify entries for operations. Expand the **Advanced** section and enter the statement in the **Custom Select** field.

What to do next

After you configure the endpoint, you can create a flow to define the relationship between the endpoint and the target directory server.

Configuring an LDAP endpoint

To configure an LDAP directory as an endpoint, you must specify the LDAP URL, login name with credentials, the search base, and root suffix.

Before you begin

Ensure that you create an endpoint and specify the type as **LDAP**. See "Configuring endpoints" on page 11.

Procedure

1. On the LDAP endpoint configuration page, in the **LDAP URL** field, enter the LDAP URL of the LDAP directory that you want to access.

The LDAP URL format is `ldap://hostname:port` or `ldap://server_IP_address:port`.

For example: `ldap://localhost:389`

Note: The default LDAP port number is 389. If you are using SSL, the default LDAP port number is 636.

2. In the **User Login** and **Password** fields, enter the distinguished name and credentials for authentication to the service.

   For example: `cn=administrator,cn=users,dc=your_domain,dc=com`

3. In the **Include entries from the following container** field, enter the search base in the LDAP directory that is polled for changes. Alternately, you can click **Contexts** and select from the **LDAP Search Base** list and then click **OK**.

   For example: `dc=your_domain,dc=com`

4. To verify the LDAP directory connection settings, click **Test Connection**. If the connection is successful, the attributes in the endpoint are displayed in a separate pane. You can use the **Filter** field to search the attributes.

5. Optional: You can also configure the following advanced parameters. Expand the **Advanced** section to view these parameters.

   Binary Attributes
   Specify a list of attributes that must be interpreted as binary values instead of strings. When you enter the attribute names in this field, enter one attribute per line and do not use any separators.

   Page Size
   Specify the number of entries per page must be returned by the request. The default value is 500.

What to do next

After you configure the endpoint, you can create a flow to define the relationship between the endpoint and the target directory server.

Configuring a Sun Directory endpoint

To configure a Sun Directory as an endpoint, you must specify the LDAP URL, login name with credentials, the search base, and root suffix.

Before you begin

Ensure that you create an endpoint and specify the type as **Sun Directory**. See "Configuring endpoints" on page 11.

Procedure

1. On the Sun Directory endpoint configuration page, in the **LDAP URL** field, enter the LDAP URL of the Sun Directory service that you want to access.

   The LDAP URL format is `ldap://hostname:port` or `ldap://server_IP_address:port`.

   For example: `ldap://localhost:389`

   Note: The default LDAP port number is 389. If you are using SSL, the default LDAP port number is 636.

2. In the **User Login** and **Password** fields, enter the distinguished name and credentials for authentication to the service.

For example: `cn=administrator,cn=users,dc=your_domain,dc=com`

3. In the **Include entries from the following container** field, enter the search base in the Sun Directory that is polled for changes. Alternately, you can click **Contexts** and select from the **LDAP Search Base** list and then click **OK**.

   For example: `dc=your_domain,dc=com`

4. To verify the Sun Directory connection settings, click **Test Connection**. If the connection is successful, the attributes in the endpoint are displayed in a separate pane. You can use the **Filter** field to search the attributes.
5. Optional: You can also configure the following advanced parameters. Expand the **Advanced** section to view these parameters.

   Seconds Before Timeout
   Specify the maximum number of seconds to wait for the next changed Sun Directory object. The default value is 0.

   Seconds Between Polling
   Specifies the number of seconds the Connector sleeps between successive polls. The default value is 60.

   Change State Key
   Specifies the name of the key or parameter that stores the change detection iterator state. The state key is used between runs to remember the last changed that was processed. If synchronization was stopped for any reason, when it is restarted, it can pick up from where it stopped.

   The value of this key must be unique for each endpoint. If you do not set this parameter, a value is computed automatically to ensure uniqueness.

   Binary Attributes
   Specify a list of attributes that must be interpreted as binary values instead of strings. When you enter the attribute names in this field, enter one attribute per line and do not use any separators.

   Page Size
   Specify the number of entries per page that must be returned by the request.

What to do next

After you configure the endpoint, you can create a flow to define the relationship between the endpoint and the target directory server.

Configuring an IBM Security Directory Server source endpoint

To configure an IBM Security Directory Server as a endpoint, you must specify the LDAP URL, login name with credentials, the search base, and root suffix.

Before you begin

Ensure that you create an endpoint and specify the type as **IBM Security Directory Server**. See “Configuring endpoints” on page 11.

Procedure

1. On the IBM Security Directory Server source endpoint configuration page, in the **LDAP URL** field, enter the LDAP URL of the IBM Security Directory Server that you want to access.

The LDAP URL format is `ldap://hostname:port` or `ldap://server_IP_address:port`.

For example: `ldap://localhost:389`

Note: The default LDAP port number is 389. If you are using SSL, the default LDAP port number is 636.

2. In the **User Login** and **Password** fields, enter the distinguished name and credentials for authentication to the server.

   For example: `cn=root`

3. In the **Include entries from the following container** field, enter the directory server search base that is polled for changes. Alternately, you can click **Contexts** and select from the **LDAP Search Base** list and then click **OK**.

   For example: `o=sample`

4. To verify the IBM Security Directory Server connection settings, click **Test Connection**. If the connection is successful, the attributes in the endpoint are displayed in a separate pane. You can use the **Filter** field to search the attributes.

5. Optional: You can also configure the following advanced parameters. Expand the **Advanced** section to view these parameters.

   Seconds Before Timeout
   Specify the maximum number of seconds to wait for the next changed directory server object. The default value is 0.

   Seconds Between Polling
   Specifies the number of seconds to sleep between successive polls. The default value is 60.

   Change State Key
   Specifies the name of the key or parameter that stores the change detection iterator state. The state key is used between runs to remember the last changed that was processed. If synchronization was stopped for any reason, when it is restarted, it can pick up from where it stopped.

   The value of this key must be unique for each endpoint. If you do not set this parameter, a value is computed automatically to ensure uniqueness.

   Binary Attributes
   Specify a list of attributes that must be interpreted as binary values instead of strings. When you enter the attribute names in this field, enter one attribute per line and do not use any separators.

   Page Size
   Specify the number of entries per page that must be returned by the request.

What to do next

After you configure the endpoint, you can create a flow to define the relationship between the endpoint and the target directory server.

Creating a flow

Create a flow that defines the relationship between the endpoints and the target IBM Security Directory Server.

Before you begin

You must create flows only after you complete the following steps:

- Connect to a target directory server.
- Configure one or more endpoints.

Procedure

1. Click the **Flows** tab to view the Flows page.
2. On the Flows page, click **Add**.
3. In the Add Flow window, specify the **Name** for the flow.
4. From the **Select endpoint type** list, select one of the configured endpoints to provide data for the flow.
5. Click **OK** to create the flow.

What to do next

After you create a flow, you can edit it to define the flow settings..

Defining flow settings

After you create a flow, you can edit the flow to define specific settings or use the default values that are provided for most settings.

Before you begin

Ensure that you create a flow.

Procedure

1. To specify or modify the flow settings, on the **Flows** page, click the name of the flow and then click **Edit**. The configuration page for the selected flow is opened. You can view and edit the flow settings in the **Source** tab.
2. To change the endpoint, from the **Source** list, select one of the configured endpoints to provide data for the flow.
3. You can specify the flow settings that are grouped into following categories:

   General settings

   Types of Entries to Handle
   Select the types of entries that must be considered for flow operations.

   By default, the options to **Handle Person entries** and **Handle Group entries** are both selected.

   Mirror the source hierarchy into Directory Server
   Specify how the hierarchy must be handled during synchronization.

   Select this check box to preserve the containers and copy the directory information tree structure from the endpoint to the target directory server.

   Clear this check box to flatten the hierarchy by pulling all entries from multiple containers in the endpoint into one specified container in the target directory.

Target container in Directory Server
Specify the search base in the target directory server.

This field is enabled only if you selected the option to mirror the source hierarchy.

This value is used as the root when mirroring the source hierarchy.

Target container for Users
Specify the container under which `Person` entries must be written.

This value is used only if you are flattening the source hierarchy.

Target container for Groups
Specify the container under which `Group` entries must be written.

This value is used only if you are flattening the source hierarchy.

Debug log output
Select this check box to generate detailed log messages with extra information, including errors about entries that were not processed or synchronized.

Filtering details
Specify the filtering criteria, one on each line in the following fields. The entries can be full DNs or partial texts.

Include the following
Specify the list of nodes in the endpoint that you want to synchronize.

The values are used for substring searches in the returned entry DNs.

Exclude the following
Specify the list of nodes in the endpoint that you want to exclude when synchronizing.

User/Person settings
Typical default values are provided for the following settings, according to type of endpoint that you selected for the flow.

Source Person Entry Object Class
Specify the object class for `Person` entries in the endpoint.

Target Person Entry Object Class
Specify the entry that must be used for creating `Person` entries in the target directory.

Source User RDN® attribute
Specify the attribute that is used as relative DN in the DN for the `Person` entries.

Target User RDN attribute
The attribute to use as the RDN for entries that are written to SDS.

Group settings

Typical default values are provided for the following settings, according to type of endpoint that you selected for the flow.

Source Group Entry Object Class

Specify the object class for `Group` entries in the endpoint.

Target Group Entry Object Class

Specify the entry that must be used for creating `Group` entries in the target directory.

Target Group Membership attribute

Specify the attribute for holding group membership in the target directory.

AssemblyLines To Call After Write Operation

You can specify the AssemblyLines that must be called after each type of write operation:

When Person Added

When Person Modified

When Person Deleted

Before you can use these fields, you must create an AssemblyLine by using the Configuration Editor. Ensure that the configuration XML file for the AssemblyLine project is copied to the `configs` folder of the Solution Directory. For more information, see the IBM Security Directory Integrator documentation and search for *Configuration Editor*.

In the **AssemblyLines To Call After Write Operation** fields, enter the following details of the AssemblyLine project that you created in Configuration Editor:

- Name of the IBM Security Directory Integrator project that contains the AssemblyLine
- Name of the AssemblyLine

Use the following format for entering the names in these fields:

Project Name`:/AssemblyLines/`*AssemblyLine Name*

For example, if your project is named `OS400` and it contains an AssemblyLine named `ReadUsers`, then you would enter:

```
OS400:/AssemblyLines/ReadUsers
```

Additional Advanced Settings

Custom property settings

You can optionally specify custom properties that are used to override the settings that are specified in the Federated Directory Server console.

Enter each custom property on a separate line.

For example: On the **General settings** page, you can enable **Debug log output** to generate detailed logs. To override this setting, enter the following custom property setting: `global.debug=true`. This setting is passed to the IBM Security Directory Integrator solution.

You can find the custom property name for a setting in the IBM Security Directory Integrator Configuration Editor:

a. On the connector configuration page, click the **Connection** tab.
b. Click the edit icon next to a field to open the Expression Editor.
c. The **Internal name** that is displayed is the custom property name.

What to do next

If you want to delete a flow that is not required, close the configuration page for that flow. On the **Flows** page, click the name of the flow, and then click **Delete Flow**. Click **OK** when the confirmation message appears.

1. Further, you can configure the following enhancements for the flow:
   - Customize attribute maps
   - Define joins
   - Configure pass-through authentication
   - Enable write-back
2. After you complete defining all of the flow settings, run the initial synchronization operation.
3. Then, either manually run incremental synchronization or schedule periodic synchronization.

Extending attribute maps for a flow

All flow relationships can contain advanced mapping and data transformation. When you set up a flow, you can specify the custom attribute maps that must be applied during the flow operations. You can choose from the attribute maps that you defined earlier for users and groups and extend those maps for a specific flow.

Before you begin

Customize attribute maps.

About this task

The custom attribute map is used to convert the attributes from the source endpoint schema to the corresponding attribute in the target schema.

Procedure

1. On the **Flows** tab, click the name of the flow and then click **Edit** to open the flow configuration page, if you did not already do so.
2. On the flow configuration page, click the **Attribute Maps** tab and then click **Person Objects** or **Group Objects** to view the custom mapping for users or groups.
3. From the **Select map for person objects** or **Select map for group objects** list, specify the map that you want to apply to the flow operations.

   The default is `person.map` for Person Objects and `group.map` for Group Objects.

   You can select another map from the list. The list includes both the ready-to-use custom attribute maps that are provided with Federated Directory Server and the maps that you customized earlier.
4. If you want to extend the attribute mapping, click **Add Attribute**.

5. In the Add Attribute window, select the attribute from the list of attributes in the target directory server. A new row is displayed with the selected attribute name under the **Directory Server Attribute** column.
6. Under **Endpoint Attribute / Assignment**, specify the attribute name in the source endpoint that must map to the target attribute.
7. Double-click the endpoint attribute name to specify more settings for the attribute mapping.
   a. Select **Enabled** to use this attribute mapping for the endpoint.
   b. Click **Simple Assignment** or **Scripted Assignment** to specify the type of mapping. If you select **Scripted Assignment**, you can define the assignment by writing JavaScript code or by calling a function in the *Solution Directory*\LDAPSync\customScript.js file. For more information, see the IBM Security Directory Integrator documentation at http://pic.dhe.ibm.com/infocenter/tivihelp/v2r1/topic/com.ibm.IBMDI.doc_7.2/welcome.html and search for *Scripting in IBM Security Directory Integrator*.
   c. Specify whether you want this mapping to be used for all operations, or only when either modifying an entry or creating an entry.
8. To delete the mapping for a specific attribute, click the check box on that row. Then, click **Remove Attribute** and click **OK** when the confirmation message appears.

Results

As a precautionary measure, when you extend the custom attribute map, the changes are made in a copy of the original attribute map file. The new file is specific to this flow. It is named with the prefix Flow_*flow_name*. For example: Flow_ADFlow_person.map.

All attribute maps are stored in the *Solution Directory*\LDAPSync directory.

Configuring a join

To augment and enrich the data from the endpoint, you can configure the flow to specify a join from another data source selectively.

About this task

A flow can join data in one endpoint with data from another endpoint. For example, a database might contain information about people, which is not available in an LDAP directory. By joining the LDAP directory with the database, Federated Directory Server can show richer data about the people.

Whenever an entry comes in from the endpoint, the flow looks it up on the join data source, merges it with the data from the endpoint, and then adds to the target IBM Security Directory Server.

Note: Only endpoints that support lookup can be used for a join. For example, endpoints like LDAP support lookup by using a certain criteria, hence they can be used for a join. File-based endpoints do not support lookup, hence cannot be used for join.

Procedure

1. On the **Flows** tab, click the name of the flow and then click **Edit** to open the flow configuration page, if you did not already do so.

2. Click the **Join** tab to view and edit the properties for the directory or data source for the join.
3. Select **Enabled** to apply the join to this flow.
4. From the **Select endpoint** list, select the endpoint that you want to use for the join. The **Select endpoint** list displays all the endpoints that you configured in Federated Directory Server. If you clear the **Enabled** check box, the **Select endpoint** field is disabled and the settings that you entered earlier are retained, but not applied during the flow operation.
5. Specify the action that must be taken when an error or failure occurs with an entry from the join during the flow operation. From the **On join failure** list, select one of the following options:
   - **Ignore error and continue** If you select this option, the error is ignored, the entry is added, modified, or deleted, and the flow operation continues with the next entry.
   - **Skip the current entry and continue** If you select this option, the entry that caused the error is skipped and the flow operation continues.
   - **Abort and terminate the flow** If you select this option, the flow operation is terminated at this entry.

   If you enabled **Debug log output** in **General Settings** on the **Source** tab, then you can view the details about the entries that caused errors.
6. You can choose to use a statement to specify simple criteria or a script for advanced criteria.
   - To specify simple criteria to find matching entries in the join, leave the **Scripted criteria** check box cleared and specify the criteria statement:
     – In the **Attribute** field, enter the attribute from the join endpoint.
     – From the **Operator** list, select the appropriate operator for the statement.
     – In the **Value** field, enter the corresponding attribute from the main endpoint.
   - To use a script to specify advanced criteria, select **Scripted criteria**. A field is provided where you can write the script for the criteria. For more information, see the IBM Security Directory Integrator documentation at http://pic.dhe.ibm.com/infocenter/tivihelp/v2r1/topic/com.ibm.IBMDI.doc_7.2/welcome.html and search for *Scripting in IBM Security Directory Integrator*.
7. Under **Attribute Maps**, you can add, remove, or modify the attribute mapping for the join.
   a. Click **Add Attribute** and select the attribute from the list of attributes in the target directory server. A new row is displayed with the selected attribute name under the **Directory Server Attribute** column.
   b. Under **Endpoint Attribute / Assignment**, specify the attribute name in the endpoint that must map to the target attribute.
   c. Double-click the endpoint attribute name to specify more settings for the attribute mapping.
      - Select **Enabled** to use this attribute mapping for the endpoint.
      - Click **Simple Assignment** or **Scripted Assignment** to specify the type of mapping. If you select **Scripted Assignment**, you can define the assignment by writing JavaScript code or by calling a function in the *Solution Directory*\LDAPSync\customScript.js file. For more information, see the IBM Security Directory Integrator documentation at

http://pic.dhe.ibm.com/infocenter/tivihelp/v2r1/topic/com.ibm.IBMDI.doc_7.2/welcome.html and search for *Scripting in IBM Security Directory Integrator*.

- Specify whether you want this mapping to be used for all operations, or only when either modifying an entry or creating an entry.

d. To delete the mapping for a specific attribute, click the check box on that row. Then, click **Remove Attribute** and click **OK** when the confirmation message appears.

8. You can also provide your own AssemblyLines instead of just endpoints to define the join operation. Expand the **Customize lookup/join assemblylines** section and specify **Person Objects**, **Group Objects**, and **Container Objects**.

   Before you can use these fields, you must create an AssemblyLine by using the Configuration Editor. Ensure that the configuration XML file for the AssemblyLine project is copied to the `configs` folder of the Solution Directory. For more information, see the IBM Security Directory Integrator documentation and search for *Configuration Editor*.

   In the **Customize lookup/join assemblylines** fields, enter the following details of the AssemblyLine project that you created in Configuration Editor:

   - Name of the IBM Security Directory Integrator project that contains the AssemblyLine
   - Name of the AssemblyLine

   Use the following format for entering the names in these fields:

   `Project Name``:/AssemblyLines/`*`AssemblyLine Name`*

   For example, if your project is named `OS400` and it contains an AssemblyLine named `ReadUsers`, then you would enter:

   ```
   OS400:/AssemblyLines/ReadUsers
   ```

Configuring pass-through authentication

You can configure a flow to delegate authentication back to the endpoint by using the pass-through authentication feature. You can use this optional feature if you want to retain the authentication credentials only in the endpoint and not in the target IBM Security Directory Server.

About this task

Pass-through authentication is a feature of IBM Security Directory Server, which delegates authentication of users to a different LDAP server. If you configure pass-through authentication for a flow, then IBM Security Directory Server attempts to verify the credentials from an external directory server on behalf of the client.

For more information about pass-through authentication, see the IBM Security Directory Server documentation at http://pic.dhe.ibm.com/infocenter/tivihelp/v2r1/topic/com.ibm.IBMDS.doc_6.3.1/welcome.htm and search for *Pass-through authentication*.

Procedure

1. On the **Flows** tab, click the name of the flow and then click **Edit** to open the flow configuration page, if you did not already do so.
2. Click the **Pass-through Authentication** tab. The four steps to enable pass-through authentication are displayed.

3. You must first verify the credentials that IBM Security Directory Server must use to access the user information in the pass-through authentication server. Specify the **Username** and **Password** and then click **Verify**.
4. Then, configure IBM Security Directory Server for pass-through authentication.
   a. Click the link, **Click to reload server PTA entry** to refresh any changes that were made in the back-end directory.
   b. Select **Enabled** to specify that the pass-through authentication mechanism must be used for this flow.
   c. The endpoint details are pre-filled based on the configuration parameters that you specified when you created the endpoint. If you require to change them you can edit the **Target subtree**, **Attribute mapping**, **Source subtree**, **Source bind DN**, and **Source bind password** fields.

   Pass-through authentication is enabled only for the users in the containers of the target subtree.
5. You must manually restart IBM Security Directory Server for the changes to take effect and to enable pass-through authentication for this flow.
6. To test that the pass-through authentication mechanism is working for this flow, specify sample user credentials **Username** and **Password** and then click **Verify**.

   You can also check the detailed logs to ensure that there are no errors in the pass-through authentication mechanism.

Enabling write-back for flows

Changes that are made in the target directory server can be propagated back to the endpoint by enabling write-back in a flow for selected attributes.

Before you begin

A global write-back option is provided as a safety feature, which you can use to turn off write-back for all flows. However, when you turn off the write-back feature globally, it prevents write-back for all flows, including the specific flows where you might want to enable write-back. Hence, you must first ensure that the write-back feature is enabled at a global level for all flows.

To enable the global write-back feature, in the navigation pane, under Directory Server, click **Write-back** and then select **Write-back enabled**. A green tick mark is displayed next to **Write-back**.

After you enable the global write-back feature, you must complete the steps in the following procedure to enable write-back for a specific flow.

About this task

Only the changes that are made to person entries that are targets of this flow are candidates for write-back operations.

Procedure

1. To enable write-back for a specific flow, on the **Flows** tab, click the name of the flow and then click **Edit**. The configuration page for the flow is opened.
2. Click the **Write-back** tab.
3. Select **Enable** to enable the write-back option for this flow.

4. Specify the attributes in the directory server that must trigger a write-back operation and map it to the attribute in the endpoint.
   a. Click **Add Attribute** and select the attribute from the list of attributes in the endpoint. A new row is displayed with the selected attribute name under the **Endpoint Attribute** column.
   b. Under **Directory Server Attribute / Assignment**, specify the attribute name in the directory server that must map to the endpoint attribute.
   c. Double-click the directory server attribute name to specify more settings for the attribute mapping.
      - Select **Enabled** to use this attribute mapping for write-back operations.
      - Click **Simple Assignment** or **Scripted Assignment** to specify the type of mapping. If you select **Scripted Assignment**, you can define the assignment by writing JavaScript code or by calling a function in the *Solution Directory*\LDAPSync\customScript.js file. For more information, see the IBM Security Directory Integrator documentation at http://pic.dhe.ibm.com/infocenter/tivihelp/v2r1/topic/com.ibm.IBMDI.doc_7.2/welcome.html and search for *Scripting in IBM Security Directory Integrator*.
   d. To delete the mapping for a specific attribute, click the check box on that row. Then, click **Remove Attribute** and click **OK** when the confirmation message appears.

Results

When a write-back operation happens, a summary of what was written back to the endpoint is displayed. The summary includes details such as the name of the flow, modified attributes, and the DNs of the directory server and endpoint is displayed. You can use the **Filter** field for searching the write-back summary.

Verifying the flow configuration

After you configure the flow and specify the criteria for the flow operations, you can run a simulated synchronization to verify the flow.

Before you begin

Ensure that you create and define a flow.

About this task

The simulated synchronization runs the same operations as an initial synchronization, but does not write anything to the directory server. This feature is helpful in the initial planning phase to verify that the flow is able to select the correct data subset in the endpoint.

Procedure

1. On the **Flows** page, click the name of the flow, and then click **Run Synchronization**.
2. In the Run Synchronization window, select **Simulate**.

Results

A complete synchronization from the source system is simulated according to the criteria specified for the flow.

A progress bar is displayed under the **Last Activity** column. The status and logs are displayed under the flow.

What to do next

If you want to stop the simulation operation that is in progress, click the name of the flow, and then click **Terminate**. Click **OK** when the confirmation message appears.

When the operation is completed, the details of the simulation such as date, operation, and modified attributes are displayed on a new tab. You can use the **Filter** field to search the table.

You can also check the status and logs to verify that the simulated synchronization was successful or to debug errors.

After you verify your flow by running a simulated synchronization, you can run the initial synchronization to migrate data to the directory server.

Synchronizing data on the target directory

After you define the flow settings you can synchronize data from the endpoint with the target IBM Security Directory Server. You can do this either manually or set up a schedule for automated synchronization at regular intervals.

Running the initial synchronization

After you define the flow settings, you can run the initial synchronization to migrate data from the endpoint to the target IBM Security Directory Server.

Before you begin

Ensure that you create and define a flow.

About this task

Initial synchronization is a one-time operation for a flow. It selects all entries in the endpoint that match the flow criteria and updates the directory server.

Procedure

1. On the **Flows** page, click the name of the flow, and then click **Run Synchronization**.
2. In the Run Synchronization window, select **Initial Synchronization**.

Results

A complete synchronization from the source system is started according to the criteria specified for the flow. Any current synchronization state data is reset.

A progress bar is displayed under the **Last Activity** column. The status and logs are displayed under the flow.

What to do next

If you want to stop a synchronization operation that is in progress, click the name of the flow, and then click **Terminate**. Click **OK** when the confirmation message appears. Terminating a flow operation leaves it in a partially synchronized state, so it must be used with caution.

When the operation is completed, you can check the status and logs to verify that the synchronization was successful or to debug errors.

After you ensure that the initial synchronization completed successfully, you can set up a schedule for synchronization at specific intervals.

Running incremental synchronization

After you run the initial synchronization, you can incrementally synchronize data on the target IBM Security Directory Server based on the changes that are made in the endpoint. You can either run a manual synchronization or set up a schedule for automated synchronization at regular intervals.

Before you begin

- Create and define a flow.
- Run the initial synchronization for the flow.

Procedure

1. On the **Flows** page, click the name of the flow, and then click **Run Synchronization**.
2. In the Run Synchronization window, select **Incremental Synchronization**.

Results

The synchronization operation is started and a progress bar is displayed under the **Last Activity** column.

What to do next

If you want to stop a synchronization operation that is in progress, click the name of the flow, and then click **Terminate**. Click **OK** when the confirmation message appears. Terminating a flow operation leaves it in a partially synchronized state, so it must be used with caution.

When the operation is completed, you can check the status and logs to verify that the synchronization was successful or to debug errors.

To automatically run the synchronization at timed intervals, you can set up a schedule for synchronization at specific intervals.

Scheduling synchronization

You can specify a schedule to automatically run the incremental synchronization operation in a flow at timed intervals.

Before you begin

- Create and define a flow.
- Run the initial synchronization for the flow.

Procedure

1. To create a schedule for a flow operation for the first time, on the **Flows** page, under the name of the flow, click **No Schedule**. To edit a schedule that is already created, click the day and time of the next scheduled operation that is displayed under the flow.
2. In the Schedule window, click **Enabled** to activate the scheduler.
3. Select the type of schedule.
   - If you select **Timer**, the synchronization runs at the intervals specified in the schedule.
   - If you select **Keepalive**, the synchronization keeps running even if a timeout value is specified in the endpoint.
4. Select the frequency for the flow operation as either **Every Month** or **Selected Month(s)**. If you choose **Selected Month(s)**, the month names are displayed and you must select one or more months.
5. Select the days on which you want to run the flow operation from the following options: **Every Day**, **Weekdays** for specify days of the week, or **Selected day(s)** to specify the days of the month.
6. Under the **Hours/Minutes/Seconds** section, enter the time of the day when you want the flow operation to start. You can also enter the wildcard * (asterisk), a comma-separated list, or a range of numbers to specify hours, minutes, and seconds.

   For example:
   - To run the synchronization at the start of each hour, enter `*` in the **Hours** field, and then enter `0` in both the **Minutes** and **Seconds** fields.
   - To run the synchronization every 15 minutes in each hour, enter `*` in the **Hours** field, `0,15,30,45` in the **Minutes** field, and `0`, in the **Seconds** field.
7. Select **Enabled**.
8. If you anticipate that a flow operation might not complete before the next operation is scheduled to start, select **Don't start if already running**. This option is useful for operations that are of a longer duration because it prevents two instances of the same operation from running simultaneously.
9. If you want to stop the flow operation when it encounters a failure, select **Terminate schedule if assemblyline fails**. For example, you can enable this option to fix errors in the log file before a failed synchronization is automatically attempted repeatedly.
10. Click **Close** to save the schedule.

Results

The day and time of the next scheduled flow operation is displayed under the flow.

What to do next

If you do not want to use the scheduler in the future, you can clear the **Enabled** check box in the Schedule window.

Viewing logs and reports

After a synchronization activity is completed, you can view the logs to verify that it was successful.

About this task

On the **Flows** page, a summary of the flow operation is displayed under each flow with the following information:

- Number of users that were added, modified, and deleted
- Number of groups that were added, modified, and deleted
- The last activity that was run on this flow
- The total number of users and groups that were processed

When you define the general settings for the flow, if you selected the **Debug log output** option, then logs are generated with detailed information for debugging.

Procedure

1. To view the detailed logs, select the operation from the **Show logs from** list. The last operation is shown by default. You can select from any of the previous logs that are listed.

   Note: To change the number of historical log files that must be stored, see "Specifying the log settings" on page 9.
2. Click one of the following sections in the log to view a detailed report:

   Summary
   Displays the following summaries:
   - Number of Person, Group, and Container entries that were processed
   - Number of errors and warnings
   - Number of entries that were skipped and not successfully written to the target directory

   Error Log
   Displays all errors and warnings. You can use the details to troubleshoot any failures in the synchronization.

   Migration Log or Sync log
   If you are viewing the logs for the initial synchronization, the migration log is displayed, otherwise the log for synchronization operation is displayed. This log contains the details of the entire flow operation.

Known issues, limitations, and workarounds

Use the problem descriptions and their solutions that are provided to resolve issues that you might encounter when you use Federated Directory Server.

Initial synchronization fails after it retrieves Page Size values

Problem

On a Windows Server 2008 R2 system, the initial synchronization fails after it retrieves the values that are set by Page Size.

This problem is specific to operations that involve Active Directory.

Description

This problem occurs in the following scenario:

- The Active Directory on a Windows Server 2008 R2 system has many users and groups, for example, 10,000 users and 10,000 groups.
- The Page Size for the Active Directory endpoint is set to 500, which is the default value.

- A flow is defined to migrate these entries to IBM Security Directory Server.

When you run the initial synchronization operation, 500 users are migrated and then an error occurs. Then, 500 groups are migrated and an error occurs. The operation is terminated with OperationNotSupportedException that is similar to the following error:

```
2013-06-12 16:37:31,250 ERROR [AssemblyLine.Flow_ADFlow1_ReadGroups_group.7]
     - [Flow_ADFlow1_ReadGroups_group/Read Groups/Default On Error]
     - [Flow_ADFlow1_ReadGroups_group/Read Groups/Default On Error]
     - javax.naming.OperationNotSupportedException: [LDAP: error code 12
     - 00002040: SvcErr: DSID-031401E7, problem 5010 (UNAVAIL_EXTENSION), data 0
]; remaining name 'ou=set1,dc=adsync,dc=tditest,dc=internal'
     - [LDAP: error code 12 - 00002040: SvcErr: DSID-031401E7,
     problem 5010 (UNAVAIL_EXTENSION), data 0
]
Stacktrace (for support):
javax.naming.OperationNotSupportedException: [LDAP: error code 12
     - 00002040: SvcErr: DSID-031401E7, problem 5010 (UNAVAIL_EXTENSION), data 0
]; remaining name 'ou=set1,dc=adsync,dc=tditest,dc=internal'
 at com.sun.jndi.ldap.LdapCtx.mapErrorCode(LdapCtx.java:3159)
 at com.sun.jndi.ldap.LdapCtx.processReturnCode(LdapCtx.java:3045)
 at com.sun.jndi.ldap.LdapCtx.processReturnCode(LdapCtx.java:2852)
 at com.sun.jndi.ldap.LdapCtx.searchAux(LdapCtx.java:1861)
 at com.sun.jndi.ldap.LdapCtx.c_search(LdapCtx.java:1784)
 at com.sun.jndi.toolkit.ctx.ComponentDirContext.p_search(ComponentDirContext.java:398)
 at com.sun.jndi.toolkit.ctx.PartialCompositeDirContext.search(PartialCompositeDirContext.java:368)
 at javax.naming.directory.InitialDirContext.search(InitialDirContext.java:287)
 at com.ibm.di.connector.LDAPConnector.getNextEntry(LDAPConnector.java:750)
 at com.ibm.di.server.AssemblyLineComponent.executeOperation(AssemblyLineComponent.java:3355)
 at com.ibm.di.server.AssemblyLineComponent.getnext(AssemblyLineComponent.java:932)
 at com.ibm.di.server.AssemblyLine.msGetNextIteratorEntry(AssemblyLine.java:3666)
 at com.ibm.di.server.AssemblyLine.executeMainStep(AssemblyLine.java:3375)
 at com.ibm.di.server.AssemblyLine.executeCycle(AssemblyLine.java:3151)
 at com.ibm.di.server.AssemblyLine.executeCycle(AssemblyLine.java:3091)
 at com.ibm.di.fc.AssemblyLineFC.executeCycle(AssemblyLineFC.java:451)
 at com.ibm.di.fc.AssemblyLineFC.perform(AssemblyLineFC.java:272)
 at sun.reflect.GeneratedMethodAccessor77.invoke(Unknown Source)
 at sun.reflect.DelegatingMethodAccessorImpl.invoke(DelegatingMethodAccessorImpl.java:55)
 at java.lang.reflect.Method.invoke(Method.java:613)
 at com.ibm.jscript.types.JavaAccessObject.call(JavaAccessObject.java:321)
 at com.ibm.jscript.types.FBSObject.call(FBSObject.java:161)
 at com.ibm.jscript.ASTTree.ASTCall.interpret(ASTCall.java:175)
 at com.ibm.jscript.ASTTree.ASTAssign.interpret(ASTAssign.java:91)
 at com.ibm.jscript.ASTTree.ASTProgram.interpret(ASTProgram.java:119)
 at com.ibm.jscript.ASTTree.ASTProgram.interpretEx(ASTProgram.java:139)
 at com.ibm.jscript.JSExpression._interpretExpression(JSExpression.java:435)
 at com.ibm.jscript.JSExpression.interpretExpression(JSExpression.java:421)
 at com.ibm.jscript.JSExpression.evaluateValue(JSExpression.java:251)
 at com.ibm.jscript.JSExpression.evaluateValue(JSExpression.java:238)
 at com.ibm.jscript.JSExpression.evaluateValue(JSExpression.java:241)
 at com.ibm.jscript.JSInterpreter.interpret(JSInterpreter.java:57)
 at com.ibm.di.script.ScriptEngine.interpret(ScriptEngine.java:940)
 at com.ibm.di.script.ScriptEngine.interpret(ScriptEngine.java:925)
 at com.ibm.di.server.ScriptComponent.add1(ScriptComponent.java:244)
 at com.ibm.di.server.ScriptComponent.add(ScriptComponent.java:210)
 at com.ibm.di.server.AssemblyLine.msExecuteNextConnector(AssemblyLine.java:3759)
 at com.ibm.di.server.AssemblyLine.executeMainStep(AssemblyLine.java:3379)
 at com.ibm.di.server.AssemblyLine.executeMainLoop(AssemblyLine.java:2988)
 at com.ibm.di.server.AssemblyLine.executeMainLoop(AssemblyLine.java:2971)
 at com.ibm.di.server.AssemblyLine.executeAL(AssemblyLine.java:2940)
 at com.ibm.di.server.AssemblyLine.run(AssemblyLine.java:1319)

2013-06-12 16:37:31,250 ERROR [AssemblyLine.Flow_ADFlow1_ReadGroups_group.7]
     - [Flow_ADFlow1_ReadGroups_group/Read Groups/Default On Error]
     - Make sure that the search base is visible in the source system,
     for example from an LDAP browser.
Also ensure that the credentials defined for the Source connection are
     authorized to see entries in this container.
***** Start dumping: ERROR *****
                    class: 'javax.naming.OperationNotSupportedException'
            connectorname: 'Read Groups'
                exception: 'javax.naming.OperationNotSupportedException:
     [LDAP: error code 12 - 00002040: SvcErr: DSID-031401E7,
     problem 5010 (UNAVAIL_EXTENSION), data 0
]; remaining name 'ou=set1,dc=adsync,dc=tditest,dc=internal''
                  message: '[LDAP: error code 12 - 00002040: SvcErr: DSID-031401E7,
     problem 5010 (UNAVAIL_EXTENSION), data 0
]'
               operation: 'get'
                  status: 'fail'
***** End dumping: ERROR *****
***** Connector parameters: Read Groups *****
    ldapUrl: ldap://9.120.98.148:389
    ldapUsername: Administrator@adsync.tditest.internal
    ldapSearchBase: ou=set1,dc=adsync,dc=tditest,dc=internal
    ldapSearchFilter: objectClass=groupofuniquenames
    ldapSearchScope: subtree
    ldapSizeLimit: 0
    ldapPageSize: 500
    jndiExtraProviderParams: null
```

Solution

Complete the following steps to work around this issue:

1. On the Windows Server 2008 R2 Active Directory, apply the following Microsoft Knowledge Base resolution that is provided at http://support.microsoft.com/kb/977180.
2. Back up your Windows registry.
3. In the following registry setting, `HKLM\System\CurrentControlSet\Services\NTDS\Parameters`, add the string value `DSA Heuristics`.
4. Set the value to `000000000001`.
5. Restart the system.

Reference

Use the reference information to know more details about the functions and components of Federated Directory Server console.

File parsers

You can select and configure the appropriate file parser from the list that is provided in the file endpoint configuration page of the Federated Directory Server console.

CBE Parser for file endpoint

Use the CBE Parser to read XML from the input stream and convert this XML to a Common Base Event (CBE) object. When the CBE Parser reads from XML, it returns all standard CBE attributes and the CBE object as attribute of the Input Map.

To access the CBE Parser configuration parameters:

1. Add a File endpoint.
2. On the File endpoint configuration page, click **Parser** and select **CBE Parser** from the list.
3. Expand the **Parser** section to view the parameters.

Parameters

Comment
Use this field to add your comments. The comment is not considered while parsing data.

Detailed Log
Select this check box to generate detailed log messages with extra information.

You can also configure the following advanced parameters. Under the **Parser** section, expand **Advanced** to view these parameters.

Character Encoding
Specify the character encoding to use for reading or writing. The default value is `UTF-8`. When the parser reads from XML, this parameter is used only if the input source does not already have encoding defined.

The CBE Parser extends the XML Parser; therefore, the same character encoding rules apply. For more information, go to the IBM Security Directory Integrator documentation at http://pic.dhe.ibm.com/infocenter/tivihelp/v2r1/topic/com.ibm.IBMDI.doc_7.2/welcome.html and search for *Character Encoding in the XML Parser*.

Validate XML
Select this check box to indicate that the parser must validate the XML with the XSD schema that is requested from the specification.

Omit XML Declaration
Select this check box to indicate that the parser must omit an XML declaration header in the output stream.

For detailed information about the CBE Parser and its input and output map attributes, go to the IBM Security Directory Integrator documentation at http://pic.dhe.ibm.com/infocenter/tivihelp/v2r1/topic/com.ibm.IBMDI.doc_7.2/welcome.html and search for *CBE Parser*.

CSV Parser for file endpoint

Use the CSV Parser to read and write data in the comma-separated values (CSV) format.

To access the CSV Parser configuration parameters:

1. Add a File endpoint.
2. On the File endpoint configuration page, click **Parser** and select **CSV Parser** from the list.
3. Expand the **Parser** section to view the parameters.

Parameters

Field Separator
Specify the character that is used to separate each column, which is typically a comma or semicolon. The default value is a semi-colon (`;`).

Sort fields
Select this check box to write header fields in alphabetical (ascending) order. The default value is `false`.

Comment
Use this field to add your comments. The comment is not considered while parsing data.

Detailed Log
Select this check box to generate detailed log messages with extra information.

You can also configure the following advanced parameters. Under the **Parser** section, expand **Advanced** to view these parameters.

Field Names
Specify the name for each column to which the parser must read or write. You can use the **Field Separator** between the field names, or specify each name on a separate line.

The order in which you specify the column names controls the order in which the columns are written to the output file.

Enable Quoting
Select this check box to output with quotation marks during a write operation. This option is selected by default.

If you clear this check box, the field is output as is, which can cause problems. When reading, quotation marks around the field are stripped if the **Enable Quoting** check box is selected. The parser is able to read quoted attributes that contain the column separator. If **Enable Quoting**

check box is cleared, the parser returns unexpected values when the input contains fields that are delimited by quotation marks.

Quote all fields
Select this check box to output all fields independently with quotation marks, if they contain quotation mark, separator, or a new line.

Write header
Select this check box to output all the field names that are separated by the column separate on the first line. This option is selected by default.

Write BOM
Select this check box to write Byte Order Marker (BOM) to the file. You must also select **Write header** for this option to take effect.

Log long lines
Specify a maximum number of bytes for a line. The line numbers of lines that are longer than this maximum number are logged.

Combine remainder in last field
Select this check box to combine all extra fields from lines that exceed the number of defined fields into a new **`Remainder`** field. The fields, and implicitly, the number of fields, are defined by **Field Names**, or its absence, the first line of the file.

Character Encoding
Specify the character encoding to use for reading or writing.

For more information, go to the IBM Security Directory Integrator documentation at http://pic.dhe.ibm.com/infocenter/tivihelp/v2r1/topic/com.ibm.IBMDI.doc_7.2/welcome.html and search for *Character Encoding conversion*.

For detailed information about the CSV Parser and its schema, go to the IBM Security Directory Integrator documentation at http://pic.dhe.ibm.com/infocenter/tivihelp/v2r1/topic/com.ibm.IBMDI.doc_7.2/welcome.html and search for *CSV Parser*.

DSMLv1 Parser for file endpoint

Use the DSMLv1 Parser to read and write XML documents. Directory Services Markup Language v1.0 (DSMLv1) enables the representation of directory structural information as an XML document. The Parser silently ignores schema entries.

To access the DSMLv1 Parser configuration parameters:

1. Add a File endpoint.
2. On the File endpoint configuration page, click **Parser** and select **DSMLv1 Parser** from the list.
3. Expand the **Parser** section to view the parameters.

Parameters

Comment
Use this field to add your comments. The comment is not considered while parsing data.

Detailed Log
Select this check box to generate log messages with detailed debug information.

You can also configure the following advanced parameters. Under the **Parser** section, expand **Advanced** to view these parameters.

DN Attribute
Specify the attribute that is used for the distinguished name DSML attribute. The default value is `$dn`.

DSML prefix
Specify the prefix that is used on XML elements to indicate that they belong to the DSML namespace. The default value is `dsml`.

DSML namespace URI
Specify the URI that identifies this namespace. The default value is `http://www.dsml.org/DSML`.

Omit XML Declaration
Select this check box to indicate that the parser must omit an XML declaration header in the output stream.

Document Validation
Select this check box to request file validation that is based on the specified DTD or schema.

Namespace Aware
Select this check box to indicate the parser must request a namespace-aware parser.

Character Encoding
Specify the character encoding to use for reading or writing. The default value is `UTF-8`.

The DSMLv1 Parser extends the Simple XML Parser; therefore, the same character encoding rules apply. For more information, go to the IBM Security Directory Integrator documentation at http://pic.dhe.ibm.com/infocenter/tivihelp/v2r1/topic/com.ibm.IBMDI.doc_7.2/welcome.html and search for *Character Encoding in the Simple XML Parser*.

For detailed information about the DSMLv1 Parser and examples of its usage, go to the IBM Security Directory Integrator documentation at http://pic.dhe.ibm.com/infocenter/tivihelp/v2r1/topic/com.ibm.IBMDI.doc_7.2/welcome.html and search for *DSMLv1 Parser*.

DSMLv2 Parser for file endpoint

Use the DSMLv2 Parser to parse and create DSMLv2 request and response messages. Directory Services Markup Language v2.0 (DSMLv2) provides a method for expressing directory queries and updates and the results of these operations as XML documents.

To access the DSMLv2 Parser configuration parameters:

1. Add a File endpoint.
2. On the File endpoint configuration page, click **Parser** and select **DSMLv2 Parser** from the list.
3. Expand the **Parser** section to view the parameters.

Parameters

Mode Specify whether the parser operates in `Server` or in `Client` mode. In `Server` mode, requests are read and responses are written. In `Client` mode, requests are written and responses are read.

Comment
Use this field to add your comments. The comment is not considered while parsing data.

Detailed Log
Select this check box to generate detailed log messages with extra information.

You can also configure the following advanced parameters. Under the **Parser** section, expand **Advanced** to view these parameters.

Character Encoding
Specify the character encoding to use for reading or writing. The default value is `UTF-8`.

The DSMLv2 Parser extends the Simple XML Parser; therefore, the same character encoding rules apply. For more information, go to the IBM Security Directory Integrator documentation at http://pic.dhe.ibm.com/infocenter/tivihelp/v2r1/topic/com.ibm.IBMDI.doc_7.2/welcome.html and search for *Character Encoding in the Simple XML Parser*.

Binary Attributes
Specify a comma delimited list of attributes that must be treated by the parser as binary attributes. A list of attributes are provided by default, which you can modify.

On Error
Specify how the server responds to failures while processing batch request elements. The valid values are `exit` and `resume`. The default value is `exit`.

Processing
Specify the value of the **`processing`** DSML attribute for batch requests. The valid values are `sequential` and `parallel`. The default value is `sequential`.

Response Order
Specify how the server orders individual responses within the batch response. The valid values are `sequential` and `unordered`. The default value is `sequential`. If you select `sequential`, the server must return a batch response in which the individual responses maintain a positional correspondence with the individual requests.

Omit XML Declaration
Select this check box to indicate that the parser must omit an XML declaration header in the output stream.

Indent Output
Select this check box to indent the output according to the depth of the statement lines. The result is cosmetic only; it has no bearing upon the semantic content of the output file.

Soap Binding
Select this check box to create SOAP DSML message. Otherwise, the DSML messages are not wrapped in SOAP.

For detailed information about the DSMLv2 Parser, its operations, attributes, and examples, go to the IBM Security Directory Integrator documentation at http://pic.dhe.ibm.com/infocenter/tivihelp/v2r1/topic/com.ibm.IBMDI.doc_7.2/welcome.html and search for *DSMLv2 Parser*.

Fixed Record Parser for file endpoint

Use the Fixed Record Parser to read and write fixed-length text records.

To access the Fixed Record Parser configuration parameters:

1. Add a File endpoint.
2. On the File endpoint configuration page, click **Parser** and select **Fixed Record Parser** from the list.
3. Expand the **Parser** section to view the parameters.

Parameters

Column Description

Specify each column description as the field name, the offset, and length, which are separated by commas. This field is a multi-line field where you must specify one column description per line.

For example:

```
field1, 1, 12
field2, 13, 4
field3, 17, 3
```

Field names are displayed during schema discovery. The offsets start at 1; invalid values such as `0` might cause an exception.

Comment

Use this field to add your comments. The comment is not considered while parsing data.

Detailed Log

Select this check box to generate detailed log messages with extra information.

You can also configure the following advanced parameters. Under the **Parser** section, expand **Advanced** to view these parameters.

Trim values

Select this check box to remove leading and trailing spaces from fields during read operations.

Character Encoding

Specify the character encoding to use for reading or writing.

For more information, go to the IBM Security Directory Integrator documentation at http://pic.dhe.ibm.com/infocenter/tivihelp/v2r1/topic/com.ibm.IBMDI.doc_7.2/welcome.html and search for *Character Encoding conversion*.

HTTP Parser for file endpoint

Use the HTTP Parser to interpret a byte stream according to the HTTP specification.

To access the HTTP Parser configuration parameters:

1. Add a File endpoint.
2. On the File endpoint configuration page, click **Parser** and select **HTTP Parser** from the list.
3. Expand the **Parser** section to view the parameters.

Parameters

Client Mode
Select this check box to indicate that the parser must operate in client HTTP response mode. If the **Client Mode** check box is cleared, the parser operates in server mode. This option is useful only if the parser is writing an output stream.

Comment
Use this field to add your comments. The comment is not considered while parsing data.

Detailed Log
Select this check box to generate detailed log messages with extra information.

You can also configure the following advanced parameters. Under the **Parser** section, expand **Advanced** to view these parameters.

Headers as Properties
Select this check box to retrieve and set the header values as properties. If this check box is cleared, the header values are read as attributes and returned as attributes.

Character Encoding
Specify the character encoding to use for reading or writing.

For more information, go to the IBM Security Directory Integrator documentation at http://pic.dhe.ibm.com/infocenter/tivihelp/v2r1/topic/com.ibm.IBMDI.doc_7.2/welcome.html and search for *Character sets/Encoding*.

For detailed information about the HTTP Parser, its schema, and header fields, go to the IBM Security Directory Integrator documentation at http://pic.dhe.ibm.com/infocenter/tivihelp/v2r1/topic/com.ibm.IBMDI.doc_7.2/welcome.html and search for *HTTP Parser*.

IdML Parser for file endpoint

Use the IdML Parser to parse the contents of an IdML (Identity Markup Language) file. It can be used for only reading IdML documents. It relies on the XML Parser for handling the IdML files and snippets.

To access the IdML Parser configuration parameters:

1. Add a File endpoint.
2. On the File endpoint configuration page, click **Parser** and select **IdML Parser** from the list.
3. Expand the **Parser** section to view the parameters.

Parameters

Comment
Use this field to add your comments. The comment is not considered while parsing data.

Detailed Log
Select this check box to generate log messages with detailed debug information.

You can also configure the following advanced parameters. Under the **Parser** section, expand **Advanced** to view these parameters.

Character Encoding
Specify the character encoding to use for reading or writing.

For detailed information about the IdML Parser and its schema, go to the IBM Security Directory Integrator documentation at http://pic.dhe.ibm.com/infocenter/tivihelp/v2r1/topic/com.ibm.IBMDI.doc_7.2/welcome.html and search for *IdML Parser*.

JSON Parser for file endpoint

Use the JSON Parser to read and write entries in the JavaScript Object Notation (JSON) format. JSON is a lightweight data-interchange format and a subset of JavaScript programming language. JSON is built with the following two structures: an ordered list of values (array) and a collection of name-value pairs (object).

To access the JSON Parser configuration parameters:

1. Add a File endpoint.
2. On the File endpoint configuration page, click **Parser** and select **JSON Parser** from the list.
3. Expand the **Parser** section to view the parameters.

Parameters

Compact Output
Select this check box to display data in compact mode. Compact mode writes JSON data on a single unformatted line and is the default mode.

Character Encoding
Specify the character encoding to be used for reading or writing data.

Comment
Use this field to add your comments. The comment is not considered while parsing data.

Detailed Log
Select this check box to generate log messages with detailed debug information.

For detailed information about the JSON Parser, its objects and attributes, and examples of its usage, go to the IBM Security Directory Integrator documentation at http://pic.dhe.ibm.com/infocenter/tivihelp/v2r1/topic/com.ibm.IBMDI.doc_7.2/welcome.html and search for *JSON Parser*.

LDIF Parser for file endpoint

Use the LDIF Parser to read and write data that is in the LDAP Data Interchange Format (LDIF). The LDIF format is used to specify a set of directory entries or a set of changes to be applied to directory entries, but not both. An LDIF file consists of a series of records that are separated by line separators.

To access the LDIF Parser configuration parameters:

1. Add a File endpoint.
2. On the File endpoint configuration page, click **Parser** and select **LDIF Parser** from the list.
3. Expand the **Parser** section to view the parameters.

Parameters

DN Attribute Name
Specify the attribute name to use for an LDIF `dn` line. The default value is `$dn`.

Version Number
Select this check box to display a version attribute in the beginning of the output (required by RFC2849). This check box is selected by default.

Comment
Use this field to add your comments. The comment is not considered while parsing data.

Detailed Log
Select this check box to generate log messages with detailed debug information.

You can also configure the following advanced parameters. Under the **Parser** section, expand **Advanced** to view these parameters.

Binary Attributes
Specify a comma delimited list of attributes that must be treated by the parser as binary attributes.

Character Encoding
Specify the character encoding to use for reading or writing. The default value is `UTF-8`.

For more information, go to the IBM Security Directory Integrator documentation at http://pic.dhe.ibm.com/infocenter/tivihelp/v2r1/topic/com.ibm.IBMDI.doc_7.2/welcome.html and search for *Character Encoding conversion*.

Note: A conforming LDIF file must always have **Character Encoding** set to `UTF-8`. **Character Encoding** is also applied for encoding or decoding BASE64 encoded strings. BASE64 encoding looks like garbled text if you do not know how to decode it.

Only Descriptive Records
Select this check box to write only descriptive records. An LDIF file might contain change records or descriptive records. A change record describes a change that is needed for an entry. It can be identified by a `changetype` line, which is the second line immediately after the `dn` line. A descriptive record describes an entry. A correct LDIF file contains either only change records or only descriptive records.

By default, this check box is not selected.

Support language tags
Select this box if you want the parser to support language tags. When information is represented in multiple languages, the server associates language tags with attribute values.

For detailed information about the LDIF Parser, go to the IBM Security Directory Integrator documentation at http://pic.dhe.ibm.com/infocenter/tivihelp/v2r1/topic/com.ibm.IBMDI.doc_7.2/welcome.html and search for *LDIF Parser*.

Line Reader Parser for file endpoint

Use the Line Reader Parser to read single lines of data from a file. The line that is read is returned in a single attribute. The attribute named `linenumber` contains the line number, starting with 1.

Use the Line Reader Parser for reading text files only and not for binary files. If you want to copy a binary file, you can use the scriptable FTP object. For more information and examples of the FTP object, go to the IBM Security Directory Integrator documentation at http://pic.dhe.ibm.com/infocenter/tivihelp/v2r1/topic/com.ibm.IBMDI.doc_7.2/welcome.html and search for *The FTP object*.

To access the Line Reader Parser configuration parameters:

1. Add a File endpoint.
2. On the File endpoint configuration page, click **Parser** and select **Line Reader Parser** from the list.
3. Expand the **Parser** section to view the parameters.

Parameters

Attribute Name
Specify the name of the attribute that contains the line of text either read or about to be written. The default value is `line`.

Comment
Use this field to add your comments. The comment is not considered while parsing data.

Detailed Log
Select this check box to generate log messages with detailed debug information.

You can also configure the following advanced parameters. Under the **Parser** section, expand **Advanced** to view these parameters.

Character Encoding
Specify the character encoding to use for reading or writing.

For more information, go to the IBM Security Directory Integrator documentation at http://pic.dhe.ibm.com/infocenter/tivihelp/v2r1/topic/com.ibm.IBMDI.doc_7.2/welcome.html and search for *Character Encoding conversion*.

Script Parser for file endpoint

Use the Script Parser to write your own parser by using JavaScript.

To access the Script Parser configuration parameters:

1. Add a File endpoint.
2. On the File endpoint configuration page, click **Parser** and select **Script Parser** from the list.
3. Expand the **Parser** section to view the parameters.

Parameters

Script Use this field to write the user-defined script to be run. A sample script is provided by default. For more information about the objects and functions that you can use in the script, go to the IBM Security Directory Integrator

documentation at http://pic.dhe.ibm.com/infocenter/tivihelp/v2r1/topic/com.ibm.IBMDI.doc_7.2/welcome.html and search for *Script Parser*.

You can also configure the following advanced parameters. Under the **Parser** section, expand **Advanced** to view these parameters.

External Files
If you want to include external script files at run time, specify them here, one file on each line. These files are run before your script.

Include Global Scripts
Select to include scripts from the Script Library.

Character Encoding
Specify the character encoding to use for reading or writing.

For more information, go to the IBM Security Directory Integrator documentation at http://pic.dhe.ibm.com/infocenter/tivihelp/v2r1/topic/com.ibm.IBMDI.doc_7.2/welcome.html and search for *Character Encoding conversion*.

Comment
Use this field to add your comments. The comment is not considered while parsing data.

Detailed Log
Select this check box to generate log messages with detailed debug information.

For detailed information about the Script Parser, its objects, methods, and schema, go to the IBM Security Directory Integrator documentation at http://pic.dhe.ibm.com/infocenter/tivihelp/v2r1/topic/com.ibm.IBMDI.doc_7.2/welcome.html and search for *Script Parser*.

Simple Parser for file endpoint

Use the Simple Parser to read and write entries that consist of attribute name and value pairs.

The entries are in the following format:

- Each line has one `attributename:value` pair.
- Multi-valued attributes use multiple lines.
- Lines with a single period mark the end of an entry.
- **`\r`** and **`\n`** in the value is an encoding of `CR` and `LF` line breaks.

To access the Simple Parser configuration parameters:

1. Add a File endpoint.
2. On the File endpoint configuration page, click **Parser** and select **Simple Parser** from the list.
3. Expand the **Parser** section to view the parameters.

Parameters

Comment
Use this field to add your comments. The comment is not considered while parsing data.

Detailed Log
Select this check box to generate log messages with detailed debug information.

You can also configure the following advanced parameters. Under the **Parser** section, expand **Advanced** to view these parameters.

Character Encoding
Specify the character encoding to use for reading or writing.

For more information, go to the IBM Security Directory Integrator documentation at http://pic.dhe.ibm.com/infocenter/tivihelp/v2r1/topic/com.ibm.IBMDI.doc_7.2/welcome.html and search for *Character Encoding conversion*.

Simple XML Parser for file endpoint

Use the Simple XML Parser to read and write XML documents. It deals with XML data that is not more than two levels deep.

The Simple XML Parser uses the Apache Xerces and Xalan libraries. The parser gives access to the XML document through a script object called `xmldom`. The `xmldom` object is an instance of the `org.w3c.dom.Document` interface. For more information about this interface, see the W3C documentation at http://www.w3schools.com or the Oracle Java™ API documentation at http://docs.oracle.com.

Note: The "XML Parser for file endpoint" on page 48 is the improved and enhanced XML Parser.

To access the Simple XML Parser configuration parameters:

1. Add a File endpoint.
2. On the File endpoint configuration page, click **Parser** and select **Simple XML Parser** from the list.
3. Expand the **Parser** section to view the parameters.

Parameters

Root Tag
Specify the root tag that encloses entries. The default value is `DocRoot`.

Entry Tag
Specify the name of the element for entries that are passed to the parser. The default value is `Entry`.

Value Tag
Specify the name of the element for attribute values that are passed to the parser. The default value is `ValueTag`.

Comment
Use this field to add your comments. The comment is not considered while parsing data.

Detailed Log
Select this check box to generate log messages with detailed debug information.

You can also configure the following advanced parameters. Under the **Parser** section, expand **Advanced** to view these parameters.

Omit XML Declaration
Select this check box to indicate that the parser must omit an XML declaration header in the output stream.

Document Validation
Select this check box to request file validation that is based on the specified DTD or schema.

Namespace Aware
Select this check box to indicate the parser must request a namespace-aware parser.

Character Encoding
Specify the character encoding to use for reading or writing. The default value is `UTF-8`.

For more information, go to the IBM Security Directory Integrator documentation at http://pic.dhe.ibm.com/infocenter/tivihelp/v2r1/topic/com.ibm.IBMDI.doc_7.2/welcome.html and search for *Character Encoding in the Simple XML Parser*.

Indent Output
Select this check box to indent the output according to the depth of the statement lines. The result is cosmetic only; it has no bearing upon the semantic content of the output file.

For detailed information about the Simple XML Parser and examples of its usage, go to the IBM Security Directory Integrator documentation at http://pic.dhe.ibm.com/infocenter/tivihelp/v2r1/topic/com.ibm.IBMDI.doc_7.2/welcome.html and search for *Simple XML Parser*.

SOAP Parser for file endpoint

Use the SOAP Parser to read and write SOAP XML documents.

The SOAP Parser converts SOAP XML documents to or from entry objects in the following manner:

- When the parser writes to the XML document, it uses attributes from the entry to build the document. The **`SOAP_CALL`** attribute is expected to contain the value for the SOAP call.
- When the parser reads from the XML document, the **`SOAP_CALL`** attribute is set to reflect the first tag that follows the `SOAP-ENV:Body` tag. For each attribute in the entry, a tag with that name and value is created. Each tag under the `SOAP_CALL` tag translates into an attribute in the entry object.

To access the SOAP Parser configuration parameters:

1. Add a File endpoint.
2. On the File endpoint configuration page, click **Parser** and select **SOAP Parser** from the list.
3. Expand the **Parser** section to view the parameters.

Parameters

Comment
Use this field to add your comments. The comment is not considered while parsing data.

Detailed Log
Select this check box to generate log messages with detailed debug information.

You can also configure the following advanced parameters. Under the **Parser** section, expand **Advanced** to view these parameters.

Omit XML Declaration
Select this check box to indicate that the parser must omit an XML declaration header in the output stream.

Document Validation
Select this check box to request file validation that is based on the specified DTD or schema.

Namespace Aware
Select this check box to indicate the parser must request a namespace-aware parser.

Character Encoding
Specify the character encoding to use for reading or writing. The default value is `UTF-8`.

For more information, go to the IBM Security Directory Integrator documentation at http://pic.dhe.ibm.com/infocenter/tivihelp/v2r1/topic/com.ibm.IBMDI.doc_7.2/welcome.html and search for *Character Encoding conversion*.

For detailed information about the SOAP Parser and examples of its usage, go to the IBM Security Directory Integrator documentation at http://pic.dhe.ibm.com/infocenter/tivihelp/v2r1/topic/com.ibm.IBMDI.doc_7.2/welcome.html and search for *SOAP Parser*.

SPMLv2 Parser for file endpoint

Use the SPMLv2 Parser to parse or write SPML Version 2 (SPMLv2) messages, which are individual SPMLv2 requests and responses.

SPMLv2 defines a core protocol over which different data models can be used to define the actual provisioning data. The combination of a data model with the SPML core specification is referred to as a profile. The use of SPML requires that a specific profile is used. This SPMLv2 Parser that is provided with Federated Directory Server console supports the SPMLv2 DSMLv2 profile.

To access the SPMLv2 Parser configuration parameters:

1. Add a File endpoint.
2. On the File endpoint configuration page, click **Parser** and select **SPMLv2 Parser** from the list.
3. Expand the **Parser** section to view the parameters.

Parameters

Comment
Use this field to add your comments. The comment is not considered while parsing data.

Detailed Log
Select this check box to generate log messages with detailed debug information.

You can also configure the following advanced parameters. Under the **Parser** section, expand **Advanced** to view these parameters.

Binary Attributes
Specify a comma delimited list of attributes that must be treated by the parser as binary attributes. A list of attributes are provided by default, which you can modify.

Character Encoding
Specify the character encoding to use for reading or writing. The default value is `UTF-8`.

The SPMLv2 Parser extends the XML Parser; therefore, the same character encoding rules apply. For more information, go to the IBM Security Directory Integrator documentation at http://pic.dhe.ibm.com/infocenter/tivihelp/v2r1/topic/com.ibm.IBMDI.doc_7.2/welcome.html and search for *Character Encoding in the XML Parser*.

For detailed information about the SPMLv2 Parser, its operations and attributes, and examples of its usage, go to the IBM Security Directory Integrator documentation at http://pic.dhe.ibm.com/infocenter/tivihelp/v2r1/topic/com.ibm.IBMDI.doc_7.2/welcome.html and search for *SPMLv2 Parser*.

XML Parser for file endpoint

Use the XML Parser to read and write XML documents. The XML Parser uses the XLXP implementation of the StAX (JSR-173) specification. StAX is a cursor-based XML Parser that can both read from and write to XML.

This XML Parser is much faster than the traditional DOM-based Simple XML Parser because it does not need to load the whole XML structure in memory like DOM does.

To access the XML Parser configuration parameters:

1. Add a File endpoint.
2. On the File endpoint configuration page, click **Parser** and select **XML Parser** from the list.
3. Expand the **Parser** section to view the parameters.

Parameters

Simple XPath
Specify the value that is used (an expression similar to XPath) to discover elements to interpret them as entries. This parameter is also used to display the structure of the XML document to be written.

Entry Tag
Specify the name of the element that holds each entry that is passed to the XML Parser.

Value Tag
Specify the name of the element that holds each attribute value that is passed to the XML Parser.

Comment
Use this field to add your comments. The comment is not considered while parsing data.

Detailed Log
Select this check box to generate log messages with detailed debug information.

You can also configure the following advanced parameters. Under the **Parser** section, expand **Advanced** to view these parameters.

Prefix to Namespace map
Specify the list of mappings between the prefix and namespace in the following format: *prefix=namespace*.

Separate each mapping with a vertical bar (|).

If the prefix starts with $, it is considered as a default namespace declaration.

The default value is `prefix=namespace`.

XSD Schema Location
Specify the schema location, which is used for display purposes only.

Character Encoding
Specify the character encoding to use for reading or writing. The default value is `UTF-8`.

For more information, go to the IBM Security Directory Integrator documentation at http://pic.dhe.ibm.com/infocenter/tivihelp/v2r1/topic/com.ibm.IBMDI.doc_7.2/welcome.html and search for *Character Encoding in the XML Parser*.

Static Attribute Declarations
Specify the declarations for attributes and prefixes. They are written with the static elements that are specified in the **Simple XPath** field.

The following text is provided in this field by default:

```
<!-- this is an example for statically declared XML attributes/namespaces -->
<!-- DocRoot xmlns="defaultNS" attr1="val2">
<Entry xmlns:p1="p1NS" p1:attr2="val2" />
</DocRoot-->
```

Ignore repeating XML declarations while reading
Select this check box to always acknowledge the first XML declaration and to ignore the subsequent declarations.

Coalescing
Select this check box to coalesce adjacent character data sections.

Omit XML declaration when writing
Select this check box to suppress writing an XML declaration to the output. This option is useful for appending to an existing XML file.

Multi-rooted Document
Select this check box to output each entry as a stand-alone element, which creates a multi-rooted document.

Indent Output
Select this check box to indent the output according to the depth of the statement lines. The result is cosmetic only; it has no bearing upon the semantic content of the output file.

Permit invalid XML characters when writing
Select this check box to include the invalid XML characters in the XML tags. If this check box is not selected, an exception occurs during write operations on the XML document.

For detailed information about the XML Parser and examples of its usage, go to the IBM Security Directory Integrator documentation at http://pic.dhe.ibm.com/infocenter/tivihelp/v2r1/topic/com.ibm.IBMDI.doc_7.2/welcome.html and search for *XML Parser*.

XML SAX Parser for file endpoint

Use the XML SAX Parser to read large XML documents that the DOM-based XML Parser cannot handle because of memory constraints. The XML SAX Parser is based on the Apache Xerces library.

The XML SAX Parser extracts data that is enclosed within the **Group tag** that you specify in the configuration. It creates an entry with the attributes that are present in the data.

To access the XML SAX Parser configuration parameters:

1. Add a File endpoint.
2. On the File endpoint configuration page, click **Parser** and select **XML SAX Parser** from the list.
3. Expand the **Parser** section to view the parameters.

Parameters

Group Tag
Specify the names of one or more XML group tags that enclose the entries. You can specify multiple tags by separating each tag name with a comma. If you do not specify a value, the root tag is used and the entire XML document is returned as a single entry.

Comment
Use this field to add your comments. The comment is not considered while parsing data.

Detailed Log
Select this check box to generate log messages with detailed debug information.

You can also configure the following advanced parameters. Under the **Parser** section, expand **Advanced** to view these parameters.

Remove prefix
Specify the prefix that you want to remove from the attribute names.

Ignore Attributes
Select this check box to ignore the attributes of the group tag and its child attributes.

Character Encoding
Specify the character encoding to use for reading or writing. The default value is `UTF-8`.

Document Validation
Select this check box to request file validation that is based on the specified DTD or schema.

Use XSD Validation
Select this check box to use XSD instead of DTD to validate the XML file.

Namespace Aware
Select this check box to indicate the parser must request a namespace-aware parser.

Read Timeout
Specify the number of seconds after which the parser stops if no data is received.

For detailed information about the XML SAX Parser and examples of its usage, go to the IBM Security Directory Integrator documentation at http://pic.dhe.ibm.com/infocenter/tivihelp/v2r1/topic/com.ibm.IBMDI.doc_7.2/welcome.html and search for *XML SAX Parser*.

XSL-Based XML Parser for file endpoint

Use the XSL-Based XML Parser to parse XML documents in any format by using the XSL that you specify. The XML documents are parsed into attribute-value pairs and stored in the entry object.

To access the XSL-Based XML Parser configuration parameters:

1. Add a File endpoint.
2. On the File endpoint configuration page, click **Parser** and select **XSL-Based XML Parser** from the list.
3. Expand the **Parser** section to view the parameters.

Parameters

Comment
Use this field to add your comments. The comment is not considered while parsing data.

Detailed Log
Select this check box to generate log messages with detailed debug information.

You can also configure the following advanced parameters. Under the **Parser** section, expand **Advanced** to view these parameters.

Character Encoding
Specify the character encoding to use for reading or writing. The default value is `UTF-8`.

The XSL-Based XML Parser extends the Simple XML Parser; therefore, the same character encoding rules apply. For more information, go to the IBM Security Directory Integrator documentation at http://pic.dhe.ibm.com/infocenter/tivihelp/v2r1/topic/com.ibm.IBMDI.doc_7.2/welcome.html and search for *Character Encoding in the Simple XML Parser*.

Indent Output
Select this check box to indent the output according to the depth of the statement lines. The result is cosmetic only; it has no bearing upon the semantic content of the output file.

Omit XML Declaration
Select this check box to indicate that the parser must omit an XML declaration header in the output stream.

Document Validation
Select this check box to request file validation that is based on the specified DTD or schema.

Namespace Aware
Select this check box to indicate the parser must request a namespace-aware parser.

To configure the input parameters, under the **Parser** section, expand **Input**.

Use Input XSL file
Select this check box to use an input XSL file. If you select this check box, the contents of the **Input XSL** field are ignored.

Input XSL File Name
Specify the path and file name of the input XSL file that contains the matching rules for transforming the user XML to the IBM Security Directory Integrator internal format.

Input XSL
Use this editable area to enter or paste the entire input XSL.

To configure the output parameters, under the **Parser** section, expand **Output**.

Use output XSL file
Select this check box to use an output XSL file. If you select this check box, the contents of the **Output XSL** field are ignored.

Output XSL File Name
Specify the path and file name of the output XSL file that has matching rules for transforming the IBM Security Directory Integrator internal format back to user XML.

Output XSL
Use this editable area to enter or paste the entire output XSL.

For detailed information about the XSL-Based XML Parser and examples of its usage, go to the IBM Security Directory Integrator documentation at http://pic.dhe.ibm.com/infocenter/tivihelp/v2r1/topic/com.ibm.IBMDI.doc_7.2/welcome.html and search for *XSL based XML Parser*.

Chapter 2. System for Cross-Domain Identity Management

The System for Cross-Domain Identity Management (SCIM) is a standard that defines schema and protocol for identity management. You can use the SCIM service that is provided in IBM Security Directory Integrator Version 7.2 with IBM Security Directory Server as the backend directory. You can also use the SCIM connector to allow IBM Security Directory Integrator solutions to read and write to servers that support the SCIM protocol.

Overview

SCIM is emerging as a standard for user and group management and is often used instead of the traditional LDAP protocol. SCIM provides the flexibility that is required for HTTP REST, cross-enterprise, and cloud application deployments. As many cloud services do not offer an LDAP interface, you can use SCIM independent of the underlying protocols.

The SCIM protocol is an application-level, REST protocol for provisioning and managing identity data on the web. The protocol supports creation, modification, retrieval, and discovery of the core identity resources, which are users and groups, and also custom resource extensions.

Features

The SCIM specification is designed to make managing user identities in cloud-based applications and services easy, fast, and inexpensive.

SCIM provides the following features:

- It builds upon experience with existing schemas and deployments.
- It places emphasis on simplicity of development and integration.
- It applies existing authentication, authorization, and privacy models.

It aims to reduce the cost and complexity of user management operations by providing a common user schema and extension model. It also binds documents to provide patterns for exchanging this schema by using standard protocols.

For more information, see the SCIM website at http://www.simplecloud.info/.

Business scenarios

The SCIM protocol is often adopted for user and group management on non-LDAP systems. New applications, both inside the enterprise and in cloud-related scenarios, can use HTTP REST to abstract away the underlying technology.

SCIM can be used successfully in the following scenarios:

- Internal deployment of new identity services with SCIM as a provisioning protocol for long-term future use.
- Internal or external cloud where LDAP is unacceptable as protocol.
- Provisioning to SaaS applications that have SCIM as the user management interface.

For more information, see the SCIM website at http://www.simplecloud.info/ and search for *SCIM scenarios*.

SCIM service in IBM Security Directory Integrator

The SCIM service in IBM Security Directory Integrator provides a SCIM interface to the IBM Security Directory Server and a SCIM connector for servers that use the SCIM protocol.

The SCIM service is built by using IBM Security Directory Integrator itself. It is actually an IBM Security Directory Integrator assembly line that acts as a server. The backend to the SCIM server must be an IBM Security Directory Server that contains the identity data. The SCIM server receives the SCIM requests and internally connects to the IBM Security Directory Server to access the data to serve the requests.

The SCIM connector implements the SCIM protocol by using JavaScript and an HTTP Client Connector.

Supported software

The SCIM service that is provided with IBM Security Directory Integrator Version 7.2 supports IBM Security Directory Server Version 6.3.1.

The SCIM service that is implemented in IBM Security Directory Integrator Version 7.2 adheres to the SCIM 1.1 specification. For more information, see the SCIM website at http://www.simplecloud.info/ and search for *specifications*.

Supported features

The SCIM service in IBM Security Directory Integrator supports most of operation of SCIM version 1.1 with appropriate attention to changes in version 2.0.

The following features are supported in the current version of the SCIM service:

- Management of users and groups with IBM Security Directory Server as the backend directory
- Schema: Enterprise user schema extension
- JSON data type
- GET/PUT/POST/DELETE requests
- PATCH: Modifying with PATCH (HTTP) request helps consumers to send only the attributes that require modification
- Pagination
- Authentication scheme: HTTP Basic
- Filtering enables consumers to use the **filter** query parameter to request a subset of resources.
- Partial resources enable consumers to use the **attributes** query parameter to specify the attributes that must be returned in resource representations
- Sorting allows consumers to specify the order in which the resources are returned.

The current version of the SCIM server does not support:

- OAuth authentication
- Bulk updates

- Automatic limitation of number of resources returned.

Note: To get the SCIM parameter **active** to work as intended, the password policy must be turned on in the IBM Security Directory Server. To turn on the password policy, set **ibm-pwdPolicy** to `true` under `cn=pwdpolicy,cn=ibmpolicies`. This setting allows SCIM to read the **ibm-pwdAccountLocked** setting from IBM Security Directory Server. For more information about setting the password policy, see the IBM Security Directory Server documentation at http://publib.boulder.ibm.com/infocenter/tivihelp/v2r1/topic/com.ibm.IBMDS.doc_6.3.1.doc/welcome.htm and search for *Setting password policy*.

Configuration files

Before you deploy the SCIM service, you must modify the configuration files to specify connection settings, user and group mapping, and schemas.

After you install IBM Security Directory Integrator Version 7.2, you can find a folder named `SCIM` in the *`tdi_install`* directory. When you create the solution directory, either manually or when the server is started, the `SCIM` folder is automatically copied to the solution directory. Alternately, you can manually copy the `SCIM` folder to your solution directory.

The `SCIM` folder contains the following set of files, including the configurations files that you can modify to configure the setup. In most cases, you might be required to update only the `SCIM.properties` file. Other files might not require any modification.

SCIM.properties

The `SCIM.properties` file contains the following server system-specific properties, including details of the backend IBM Security Directory Server.

LDAPServer
The URL for the IBM Security Directory Server that stores the user data.

userSearchBase
The Search Base for users in the IBM Security Directory Server.

groupSearchBase
The Search Base for groups in the IBM Security Directory Server.

userObjectClass
The list of object classes that are used when a user is created in the IBM Security Directory Server.

groupObjectClass
The list of object classes that are used when a group is created in the IBM Security Directory Server.

userSearchFilter
Used to find all users in the userSearchBase.

groupSearchFilter
Used to find all groups in the groupSearchBase.

dummyGroupMember
When new groups are created, if **dummyGroupMember** has a value and there are no members in the group, this value is added to avoid object violation error.

audit.log
Set this parameter to `true` to create audit logs.

audit.logFile
The name of the audit log file.

audit.logFileDatePattern
The date pattern specifies how often the log file is rolled over to a backup file. It also specifies how the date is appended to the log file name for the backup files that store previous logs.

Location
The externally accessible URL of the SCIM service. It affects only the location headers in SCIM replies.

httpPort
The port that the SCIM Service uses for listening. The SCIM Service always uses SSL.

AuthenticationRealm
The realm that is presented to the user when asked for authentication.

audit.syslog
Indicates whether syslogging to QRadar® is enabled. Set the value to `true` to enable.

audit.QRadarHost
The host where QRadar is located.

audit.QRadarPort
The port number for QRadar.

audit.facility
The facility for the audit messages.

audit.eventID
The event ID to use in audit logs.

audit.devTimeFormat
The date format to use in audit logs.

LDAP.LookupLimit
The maximum number of resources that can be found by the SCIM Service. The default value is only 20000, to avoid memory overflow.

UserMapping.json and GroupMapping.json

The `UserMapping.json` and `GroupMapping.json` files specify the mapping between SCIM attributes and IBM Security Directory Server user or group attributes. Each entry in these files contains an SCIM attribute name and an LDAP attribute name. The entry might also contain the following extra attributes.

ReadOnly
Specifies that the value is mapped only from LDAP to SCIM and not the other way.

WriteOnly
Specifies that the value is mapped only from SCIM to LDAP and not the other way. This entry must be used for password.

CreateDN
Specifies that the value is also used to create a distinguished name (DN) in the IBM Security Directory Server, by appending the userSearchBase to the value. To be able to create new resources, there must be one entry with the **CreateDN** attribute, which uses a SCIM attribute name that is always provided.

Type
Provides the canonical type for a multi-valued attribute.

Conversion
Specifies a conversion of the attribute value. The conversion attribute can have one of the following values:

- **DateTime** converts the value from LDAP date format to SCIM date format.
- **Group** converts the value from an LDAP group to a SCIM group.
- **NewLines** converts the new lines in SCIM values to $ in LDAP values and vice versa.

Note:

- There must be only one map entry for each SCIM name, unless the entries have a unique **Type**.
- There must be only one entry for each LDAP name, unless the entries are **ReadOnly**.

UserSchema.json and GroupSchema.json

The `UserSchema.json` and `GroupSchema.json` files provide the schema definition of users or groups as per the SCIM specification. The attributes that are specified must match the attributes that are defined in the `UserMapping.json` and `GroupMapping.json` files.

ServiceProviderConfig.json
Defines the specification compliance, supported data models, authentication schemes, and so forth.

SCIM.xml
The configuration file that implements the SCIM service.

QRadarLogging.map

The `QRadarLogging.map` file specifies the values for attributes that are sent to the QRadar system when QRadar syslogging is enabled.

For more information, see the `Readme.txt` file in the `SCIM` folder in the *solution_directory* of IBM Security Directory Integrator installation.

Starting the SCIM service

Use the **ibmdisrv** command to start the SCIM service.

Before you begin

- Modify the SCIM configuration files as required.

Procedure

Run the following command:

```
ibmdisrv -c SCIM/SCIM.xml -r SCIM_Service -w
```

Results

When the SCIM service is started, it tries to do an anonymous bind to the IBM Security Directory Server. If this fails, the SCIM service stops and shows a message in the `ibmdi.log` file: CTGDIS1930E Cannot connect to the LDAP server.

SCIM connector

You can use the SCIM connector in IBM Security Directory Integrator Version 7.2 to read and write to servers that support the SCIM protocol.

The SCIM connector works like other IBM Security Directory Integrator connectors, but behind the scenes it passes REST calls and uses SCIM operations.

For information about how to configure and use the SCIM connector, see IBM Security Directory Integrator documentation at http://pic.dhe.ibm.com/infocenter/tivihelp/v2r1/topic/com.ibm.IBMDI.doc_7.2/welcome.html and search for *SCIM Connector*.

Logging and tracing

The logging and tracing feature of SCIM can help you to help find the cause of issues and resolve them.

You can set the **debug** parameter in the `SCIM.properties` file to `true` to increase the amount of data that is logged in the log files.

To configure audit logging, you can set the following properties in the `SCIM.properties` file.

audit.log
Indicates whether logging is turned on. Set the value to `true` to turn on the logging.

auditLogFile
Specifies the file name where the daily logging is done.

audit.logFileDatePattern
Specifies how often the log file must be rolled over to a new file. The default value is daily. The rollover happens only when the first message is logged in the new day. The logging is done by using a log4j DailyRollingFileAppender.

The logging is done in JSON format, where each line is one JSON object as shown in the following example:

```
{"url":"\/Users", "date":"2013-08-03 14:19:25,234","host":"127.0.0.1",
     "method":"POST","user":"cn=root",
     "resourceID":"cn=John Doe,ou=People,DC=EXAMPLE,DC=COM",
     "date":"2013-08-03 14:19:25,296","user":"cn=root","status":"201 Created"}
```

The JSON objects have the following attributes:

user
The user name that authorizes the request.

date
The date and time when the request was received.

remoteHost
The IP address of the host from which the request was received.

remotePort
The port from which the request came.

localHost
The local IP address.

localPort
The local port.

method
The method in the request.

url
The URL in the request.

userAgent
Name of the browser from which the request came, if available.

resourceID
The resource ID that was created or returned by the request.

status
The HTTP status that was returned.

SCIM object model

SCIM is built on an object model where a *Resource* is the common denominator and all SCIM objects are derived from it.

SCIM currently has three objects that directly inherit from the Resource object. The `ServiceProviderConfiguration` and `Schema` are used for discovery and contain no user information. The `CoreResource` object contains the user and group data within its two child resources, User and Group.

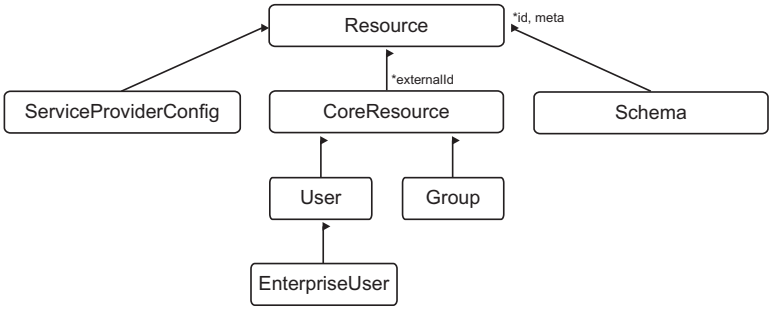

Figure 2. SCIM object model

Operations

SCIM provides a REST API with a rich but simple set of operations that you can use to manage resources.

The SCIM operations support everything from patching a specific attribute on a specific user to doing massive bulk updates.

Create `POST https://example.com/{v}/{resource}`

Read `GET https://example.com/{v}/{resource}/{id}`

Replace
`PUT https://example.com/{v}/{resource}/{id}`

Delete `DELETE https://example.com/{v}/{resource}/{id}`

Update
`PATCH https://example.com/{v}/{resource}/{id}`

Search

```
GET https://example.com/{v}/{resource}?filter={attribute}{op}{value}
     &sortBy={attributeName}&sortOrder={ascending|descending}
```

Bulk `POST https://example.com/{v}/Bulk`

Discovery operations

To simplify interoperability, SCIM provides two end points to discover supported features and specific attribute details.

`GET /ServiceProviderConfigs`
Discovers specification compliance, authentication schemes, data models.

`GET /Schemas`

- `GET /Schemas/User`
- `GET /Schemas/Group`

Introspects resources and attribute extensions.

Examples of SCIM operations

You can use the SCIM operations to search, create, modify, or delete users and groups in various scenarios.

Example 1

To get a list of all users, send the following request:

```
GET /users
```

Example 2

The following example shows how to get a list of all users but include only the **`displayName`** and **`id`** attributes. It also limits the result to the users from numbers 11 - 20.

Request:

```
GET /users?attributes=displayName,id&count=10&startIndex=11
```

Results:

```
{
  "schemas":  [
    "urn:scim:schemas:core:1.0"
  ]
,
  "Resources":  [
        {
      "id":"7b401115-35f2-4a74-8384-a684cb4f31a1",
      "displayName":"Alexander Shelton"
    }
,
        {
      "id":"44216fbe-36a1-4215-b6f7-032775bc5e07",
      "displayName":"Andy Walker"
    }
,
        {
      "id":"c5292b7e-ffeb-4855-a086-7289d3445bd6",
      "displayName":"Alan White"
    }
,
        {
      "id":"5ad2d53c-9844-48ca-8460-c0d80fec5972",
      "displayName":"Alan Worrell"
    }
,
```

```
        {
      "id":"2b62e6a0-a698-4ffb-a107-1078b2d56437",
      "displayName":"Barbara Francis"
    }
,
        {
      "id":"3904d440-3f54-46cf-b63a-aacab03ac767",
      "displayName":"Bjorn Free"
    }
,
        {
      "id":"abb9526e-dfa8-452a-9d88-9eff3d79da90",
      "displayName":"Barbara Hall"
    }
,
        {
      "id":"d7df93df-d0bd-4c60-ad52-ec2bf8917fbc",
      "displayName":"Benjamin Hall"
    }
,
        {
      "id":"f98c9470-d7fe-490f-ab71-e84c9d3e9448",
      "displayName":"Barbara Jablonski"
    }
,
        {
      "id":"87fd1385-7d13-4423-851a-fb1d047bc2f0",
      "displayName":"Bjorn Jensen"
    }

  ]
,
  "totalResults":"163",
  "startIndex":"11",
  "itemsPerPage":"10"
}
```

Example 3

The following example gets a list of all users where the **familyName** starts with k.

Request:

```
GET /users?filter=name.familyName sw "k"
```

Results:

```
{
  "schemas":  [
    "urn:scim:schemas:core:1.0"
  ]
,
  "Resources":  [
        {
      "id":"6f0fa17b-d988-4f95-98c0-095a545cc44e",
      "externalID":"aknutson",
      "meta":      {
        "created":"2013-04-16T09:14:02Z",
        "modified":"2013-04-16T09:14:02Z"
      }
,
      "userName":"uid=aknutson,ou=People,DC=EXAMPLE,DC=COM",
      "displayName":"Ashley Knutson",
      "name":      {
        "givenName":"Ashley",
        "familyName":"Knutson"
      }
```

```
,
      "phoneNumbers":      [
                {
           "type":"work",
           "value":"+1 408 555 2169"
         }
,
                {
           "type":"fax",
           "value":"+1 408 555 4774"
         }

      ]
,
      "emails":      [
                {
           "type":"work",
           "value":"aknutson@example.com"
         }

      ]

    }
,
        {
      "id":"6f7a3e28-db6c-4846-ae78-2346f39f65ee",
      "externalID":"ekohler",
      "meta":      {
        "created":"2013-04-16T09:14:02Z",
        "modified":"2013-04-16T09:14:02Z"
      }
,
      "userName":"uid=ekohler,ou=People,DC=EXAMPLE,DC=COM",
      "displayName":"Elba Kohler",
      "name":      {
        "givenName":"Elba",
        "familyName":"Kohler"
      }
,
      "phoneNumbers":      [
                {
           "type":"work",
           "value":"+1 408 555 1926"
         }
,
                {
           "type":"fax",
           "value":"+1 408 555 9332"
         }

      ]
,
      "emails":      [
                {
           "type":"work",
           "value":"ekohler@example.com"
         }

      ]

    }
,
        {
      "id":"e5318e13-1534-4eb9-9237-e1367a2744e1",
      "externalID":"skellehe",
      "meta":      {
        "created":"2013-04-16T09:14:02Z",
```

```
      "modified":"2013-04-16T09:14:02Z"
    }
,
    "userName":"uid=skellehe,ou=People,DC=EXAMPLE,DC=COM",
    "displayName":"Sue Kelleher",
    "name":      {
      "givenName":"Sue",
      "familyName":"Kelleher"
    }
,
    "phoneNumbers":      [
             {
        "type":"work",
        "value":"+1 408 555 3480"
      }
,
             {
        "type":"fax",
        "value":"+1 408 555 8721"
      }

    ]
,
    "emails":      [
             {
        "type":"work",
        "value":"skellehe@example.com"
      }

    ]

  }
,
      {
    "id":"3bac3d16-33ee-4a39-a6d1-063c5537530a",
    "externalID":"tkelly",
    "meta":      {
      "created":"2013-04-16T09:14:02Z",
      "modified":"2013-04-16T09:14:02Z"
    }
,
    "userName":"uid=tkelly,ou=People,DC=EXAMPLE,DC=COM",
    "displayName":"Timothy Kelly",
    "name":      {
      "givenName":"Timothy",
      "familyName":"Kelly"
    }
,
    "phoneNumbers":      [
             {
        "type":"work",
        "value":"+1 408 555 4295"
      }
,
             {
        "type":"fax",
        "value":"+1 408 555 1992"
      }

    ]
,
    "emails":      [
             {
        "type":"work",
        "value":"tkelly@example.com"
      }
```

```
      ]

    }

  ]
,
  "totalResults":"4"
}
```

Example 4

The following example shows how to search for the user with the **id** 2064f364-260b-4c29-8c28-b12583486ca3.

Request:

```
GET /users/2064f364-260b-4c29-8c28-b12583486ca3
```

Results:

```
{
  "id":"2064f364-260b-4c29-8c28-b12583486ca3",
  "externalID":"abergin",
  "meta":  {
    "created":"2013-04-16T09:14:02Z",
    "modified":"2013-04-16T09:14:02Z"
  }
,
  "userName":"uid=abergin,ou=People,DC=EXAMPLE,DC=COM",
  "displayName":"Andy Bergin",
  "name":  {
    "givenName":"Andy",
    "familyName":"Bergin"
  }
,
  "phoneNumbers":  [
        {
      "type":"work",
      "value":"+1 408 555 8585"
    }
,
        {
      "type":"fax",
      "value":"+1 408 555 7472"
    }

  ]
,
  "emails":  [
        {
      "type":"work",
      "value":"abergin@example.com"
    }

  ]
,
  "groups":  [
        {
      "value":"57a96228-48a6-4f29-a8ad-345828fccd6a",
      "display":"QA Managers"
    }

  ]
,
  "schemas":  [
```

```
    "urn:scim:schemas:core:1.0"
  ]

}
```

Example 5

The following example shows how to get a list of all users created after a specified date.

Request:

```
GET /users?filter=meta.created gt "2013-05-17T00:00:00Z"
```

Results:

```
{
  "schemas":  [
    "urn:scim:schemas:core:1.0"
  ]
,
  "Resources":  [
        {
      "id":"78a13de7-0ef9-42ae-ba7c-b9c64a2050aa",
      "externalID":"wlutz2",
      "meta":      {
        "created":"2013-05-21T11:39:48Z",
        "modified":"2013-05-21T11:53:30Z"
      }
,
      "userName":"uid=wlutz2,ou=People,DC=EXAMPLE,DC=COM",
      "displayName":"Wendy Lutz",
      "name":      {
        "givenName":"Wendy",
        "familyName":"Lutz"
      }
,
      "phoneNumbers":      [
                {
          "type":"work",
          "value":"+1 408 555 3358"
        }
,
                {
          "type":"fax",
          "value":"+1 408 555 9332"
        }

      ]
,
      "emails":      [
                {
          "type":"work",
          "value":"wlutz@example.com"
        }

      ]

    }
,
        {
      "id":"a4cc7512-1530-4adc-952b-cd752aa79828",
      "externalID":"wlutz4",
      "meta":      {
        "created":"2013-05-21T11:54:12Z",
        "modified":"2013-05-21T11:54:12Z"
      }
```

```
,
      "userName":"uid=wlutz4,ou=People,DC=EXAMPLE,DC=COM",
      "displayName":"Wendy Lutz",
      "name":      {
        "givenName":"Wendy",
        "familyName":"Lutz"
      }
,
      "phoneNumbers":      [
              {
          "type":"work",
          "value":"+1 408 555 3358"
        }
,
              {
          "type":"fax",
          "value":"+1 408 555 9332"
        }

      ]
,
      "emails":      [
              {
          "type":"work",
          "value":"wlutz@example.com"
        }

      ]

    }
,
        {
      "id":"9be8c033-cf93-448e-a96b-d1290ff6d445",
      "externalID":"abergin2",
      "meta":      {
        "created":"2013-05-24T11:29:51Z",
        "modified":"2013-05-24T11:51:09Z"
      }
,
      "userName":"uid=abergin2,ou=People,DC=EXAMPLE,DC=COM",
      "displayName":"Andy Bergin Jr",
      "name":      {
        "givenName":"Andy",
        "familyName":"Bergin"
      }
,
      "phoneNumbers":      [
              {
          "type":"work",
          "value":"+1 408 555 8585"
        }
,
              {
          "type":"fax",
          "value":"+1 408 555 7472"
        }

      ]
,
      "emails":      [
              {
          "type":"work",
          "value":"abergin@example.com"
        }

      ]
```

```
    }

  ]
,
  "totalResults":"3"
}
```

Example 6

To create a user, send the following request:

```
POST /users
```

The body must contain information about the new user in JSON format as shown in the following example:

```
{
      "externalID":"abergin2",
      "displayName":"Andy Bergin",
      "name":      {
        "givenName":"Andy",
        "familyName":"Bergin"
      }
,
      "phoneNumbers":      [
                {
           "type":"work",
           "value":"+1 408 555 8585"
         }
,
                {
           "type":"fax",
           "value":"+1 408 555 7472"
         }

      ]
,
      "emails":      [
                {
           "type":"work",
           "value":"abergin@example.com"
         }

      ]
}
```

Results:

```
200 OK
{
  "id":"9be8c033-cf93-448e-a96b-d1290ff6d445",
  "externalID":"abergin2",
  "meta":  {
    "created":"2013-05-24T11:29:51Z",
    "modified":"2013-05-24T11:51:09Z"
  }
,
  "userName":"uid=abergin2,ou=People,DC=EXAMPLE,DC=COM",
  "displayName":"Andy Bergin",
  "name":  {
    "givenName":"Andy",
    "familyName":"Bergin"
  }
,
  "phoneNumbers":  [
        {
      "type":"work",
```

```
      "value":"+1 408 555 8585"
    }
,
        {
      "type":"fax",
      "value":"+1 408 555 7472"
    }

  ]
,
  "emails":  [
        {
      "type":"work",
      "value":"abergin@example.com"
    }

  ]
,
  "schemas":  [
    "urn:scim:schemas:core:1.0"
  ]

}
```

Example 7

The following example shows how to modify a user. It changes only the **displayName** of the user that was created in the previous example with id b9be8c033-cf93-448e-a96b-d1290ff6d445.

Request:

```
PATCH /users/b9be8c033-cf93-448e-a96b-d1290ff6d445
```

The HTTP body must contain the following information:

```
{
 "displayName":"Andy Bergin Jr"
}
```

Results:

```
{
  "id":"9be8c033-cf93-448e-a96b-d1290ff6d445",
  "externalID":"abergin2",
  "meta":  {
    "created":"2013-05-24T11:29:51Z",
    "modified":"2013-05-24T11:51:09Z"
  }
,
  "userName":"uid=abergin2,ou=People,DC=EXAMPLE,DC=COM",
  "displayName":"Andy Bergin Jr",
  "name":  {
    "givenName":"Andy",
    "familyName":"Bergin"
  }
,
  "phoneNumbers":  [
        {
      "type":"work",
      "value":"+1 408 555 8585"
    }
,
        {
      "type":"fax",
      "value":"+1 408 555 7472"
    }
```

```
  ]
,
  "emails":  [
        {
      "type":"work",
      "value":"abergin@example.com"
    }

  ]
,
  "schemas":  [
    "urn:scim:schemas:core:1.0"
  ]

}
```

Note: To test the operations with a browser that does not have a **PATCH** command, you can set the value of the HTTP header X-HTTP-Method-Override to PATCH. You can also use this setting to work around firewalls that block certain HTTP methods.

Example 8

The following example shows how to delete the user with **id** 2064f364-260b-4c29-8c28-b12583486ca3.

Request:

```
DELETE /users/2064f364-260b-4c29-8c28-b12583486ca3
```

Results:

```
200 OK
```

Example 9

To get a list of all groups, use the following request:

```
GET /groups
```

Example 10

The following example shows how to search for a specific group by its **id**.

Request:

```
GET /groups/5653c887-1d5a-42cf-a470-6a2fe2608730
```

Results:

```
{
  "id":"5653c887-1d5a-42cf-a470-6a2fe2608730",
  "externalID":"Accounting Managers",
  "meta":  {
    "created":"2013-04-16T09:10:45Z",
    "modified":"2013-04-16T09:10:45Z"
  }
,
  "displayName":"Accounting Managers",
  "members":  [
        {
      "value":"71e064d4-3791-4ac8-b7c6-62686ce710cd",
      "display":"Sam Carter"
    }
```

```
,
        {
      "value":"6ba0ff5b-98b4-41c8-be28-331b99d94bde",
      "display":"Ted Morris"
    }

  ]
,
  "schemas":  [
    "urn:scim:schemas:core:1.0"
  ]

}
```

Example 11

The following example shows how to search for a group by its **displayName**.

Request:

```
 GET /groups?filter=displayName eq "Accounting Managers"
```

Results:

```
{
  "id":"5653c887-1d5a-42cf-a470-6a2fe2608730",
  "externalID":"Accounting Managers",
  "meta":  {
     "created":"2013-04-16T09:10:45Z",
    "modified":"2013-04-16T09:10:45Z"
  }
,
  "displayName":"Accounting Managers",
  "members":  [
        {
      "value":"71e064d4-3791-4ac8-b7c6-62686ce710cd",
      "display":"Sam Carter"
    }
,
        {
      "value":"6ba0ff5b-98b4-41c8-be28-331b99d94bde",
      "display":"Ted Morris"
    }

  ]
,
  "schemas":  [
    "urn:scim:schemas:core:1.0"
  ]

}
```

Example 12

The following example shows how to create a group.

Request:

```
POST /groups
```

The body must contain the information about the new group:

```
{
       "externalID":"Test Group",
       "displayName":"Test Group",
      "members":      [
```

```
     "5156d423-3c74-415b-844f-606a2aabafcc",
     "900faa78-d7c6-421c-9181-313134d17dd0"
      ]
    }
```

Results:

```
201 Created
{
  "id":"7e15ce9e-2fe7-4624-b5d5-adedc242e07a",
  "externalID":"Test Group",
  "meta":  {
    "created":"2013-05-27T02:37:38Z",
    "modified":"2013-05-27T02:37:38Z"
  }
,
  "displayName":"Test Group",
  "members":  [
        {
      "value":"5156d423-3c74-415b-844f-606a2aabafcc",
      "display":"Kirsten Vaughan"
    }
,
        {
      "value":"900faa78-d7c6-421c-9181-313134d17dd0",
      "display":"Robert Daugherty"
    }

  ]
,
  "schemas":  [
    "urn:scim:schemas:core:1.0"
  ]

}
```

Notices

This information was developed for products and services offered in the U.S.A. IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785 U.S.A.

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan, Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law :

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement might not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
2Z4A/101
11400 Burnet Road
Austin, TX 78758 U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to

IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

© (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. _enter the year or years_. All rights reserved.

If you are viewing this information in softcopy form, the photographs and color illustrations might not be displayed.

Trademarks

IBM, the IBM logo, and ibm.com® are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

Adobe, Acrobat, PostScript and all Adobe-based trademarks are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, other countries, or both.

IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency which is now part of the Office of Government Commerce.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.

Linear Tape-Open, LTO, the LTO Logo, Ultrium, and the Ultrium logo are trademarks of HP, IBM Corp. and Quantum in the U.S. and other countries.

Index

P

R

S

T

U

V

W

Printed in USA

SC27-6211-00