FORGEROCK®

Connector Developer's Guide

/ ForgeRock Identity Management 6.0

Latest update: 6.0.0.7

Lana Frost
László Hordós
Mark Craig

Copyright © 2012-2017 ForgeRock AS.

Abstract

Hands-on guide to developing connectors using the Open Identity Connector Framework (OpenICF). OpenICF provides connectors for a consistent generic layer between applications and target resources.

Table of Contents

Preface

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

The ForgeRock Common REST API works across the platform to provide common ways to access web resources and collections of resources.

1. About This Guide

This guide shows you how to use and develop OpenICF connectors, which decouple applications from data resources.

This guide is written for Java, C#, and web developers who use OpenICF to connect to resources from their applications, and who build their own OpenICF connectors and connector servers.

2. Accessing Documentation Online

ForgeRock publishes comprehensive documentation online:

• The ForgeRock Knowledge Base offers a large and increasing number of up-to-date, practical articles that help you deploy and manage ForgeRock software.

  While many articles are visible to community members, ForgeRock customers have access to much more, including advanced information for customers using ForgeRock software in a mission-critical capacity.

• ForgeRock product documentation, such as this document, aims to be technically accurate and complete with respect to the software documented. It is visible to everyone and covers all product features and examples of how to use them.

3. Using the ForgeRock.org Site

The ForgeRock.org site has links to source code for ForgeRock open source software, as well as links to the ForgeRock forums and technical blogs.

If you are a *ForgeRock customer*, raise a support ticket instead of using the forums. ForgeRock support professionals will get in touch to help you.

Chapter 1

About the OpenICF Framework and ICF Connectors

The Open Identity Connector Framework (OpenICF) project provides interoperability between identity, compliance and risk management solutions. An OpenICF Connector enables provisioning software, such as IDM, to manage the identities that are maintained by a specific identity provider.

OpenICF connectors provide a consistent layer between identity applications and target resources, and expose a set of operations for the complete lifecycle of an identity. The connectors provide a way to decouple applications from the target resources to which data is provisioned.

OpenICF focuses on provisioning and identity management, but also provides general purpose capabilities, including authentication, create, read, update, delete, search, scripting, and synchronization operations. Connector bundles rely on the OpenICF Framework, but applications remain completely separate from the connector bundles. This enables you to change and update connectors without changing your application or its dependencies.

Many connectors have been built within the OpenICF framework, and are maintained and supported by ForgeRock and by the OpenICF community. However, you can also develop your own OpenICF connector, to address a requirement that is not covered by one of the existing connectors. In addition, OpenICF provides two *scripted connector toolkits*, that enable you to write your own connectors based on Groovy or PowerShell scripts.

The OpenICF framework can use IDM, Sun Identity Manager, and Oracle Waveset connectors (version 1.1) and can use ConnID connectors up to version 1.4.

This guide provides the following information:

- An overview of the OpenICF framework and its components

- Information on how to use the OpenICF existing connectors in your application (both locally and remotely)

- Information on how to write your own Java and .NET connectors, scripted Groovy connectors, or scripted PowerShell connectors

1.1. Overview of the OpenICF Architecture

OpenICF is situated between the identity management application and the target resource. The framework provides a generic layer between the application and the connector bundle that accesses

the resource. The framework implements an API, that includes a defined set of operations. When you are building a connector, you implement the Service Provider Interface (SPI), and include only those operations that are supported by your target resource. Each connector implements a set of SPI operations. The API operations call the SPI operations that you implement.

The following image shows a high-level overview of an OpenICF deployment.
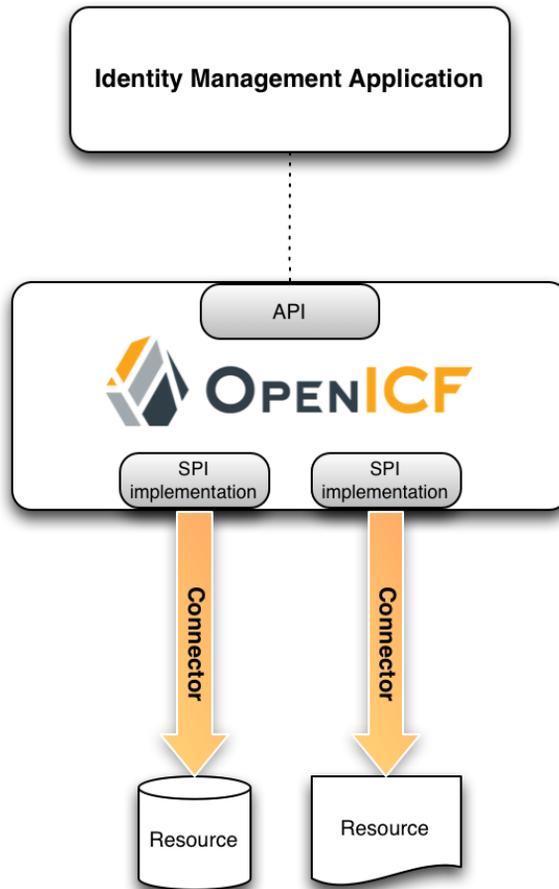
1.1.1. Understanding the OpenICF Framework Components

When you are building, or modifying, an identity management application to use the OpenICF Framework and its connectors, you use the following interfaces of the API:

• Connector Info Manager Component

The connector info manager maintains a set of connector info instances, each of which describes an available connector. The OpenICF Framework provides three different types of connector info manager:

• Local

  A local connector info manager accesses the connector bundle or assembly directly.

• Remote

  A remote connector info manager accesses the connector bundle or assembly through a remote connector server.

• OSGi

  An OSGi connector info manager accesses the connector bundle within the OSGi context.

For more information, see "Creating the Connector Info Manager".

• Connector Info Component

The connector info component provides meta information (display name, category, messages, and so forth) for a given connector.

• Connector Key Component

The connector key component uniquely identifies a specific connector implementation.

• API Configuration

The API configuration holds the available configuration properties and values from both the API, and the SPI, based on the connector type.

For more information, see "Implementing the Configuration Interface".

• Connector Facade Interface

The connector facade is the main interface through which an application invokes connector operations. The connector facade represents a specific connector instance, that has been configured in a specific way. For more information, see "Creating the Connector Facade".

When you are building a new connector, you implement the SPI, including the following interfaces:

• The `connector` interface.

The connector interface handles initialization and disposal of the connector, and determines whether the connector is poolable. For more information, see "Implementing the Connector Interface".

• The `configuration` interface.

The configuration interface implementation includes all of the required information to enable the connector to connect to the target system, and to perform its operations. The configuration interface implements getters and setters for each of its defined properties. It also provides a validate method that determines whether all the required properties are available, and valid. For more information, see "Implementing the Configuration Interface".

The OpenICF framework uses the configuration interface implementation to build the *configuration properties* inside the API configuration.

When the configuration interface is implemented, it becomes available to the default API configuration.

- Any `operations` that the target resource can support, such as `CreateOp`, `UpdateOp`, `DeleteOp` and so forth. For more information, see "Implementing the Operation Interfaces".

1.2. Overview of a Remote Connector Implementation

Connectors can run locally (on the same host as your application) or remotely (on a host that is remote to your application). Connectors that run remotely require a *connector server*, running on the same host as the connector. Applications access the connector implementation *through* the connector server.

> **Note**
>
> The OpenICF framework can support both local and remote connector implementations simultaneously.

Connector servers also enable you to run connector bundles that are written in C# on a .NET platform, and to access them over the network from a Java or .NET application.

The following image shows a high-level overview of an OpenICF deployment, including a remote connector server.
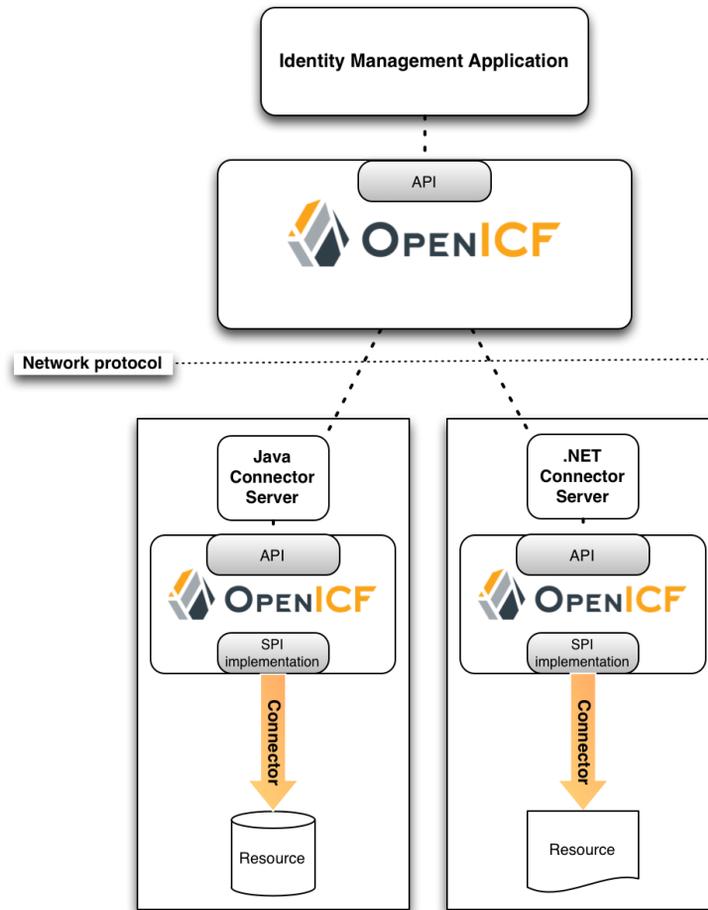
For more information about connector servers, and how to use them in your application, see "*Using OpenICF Connector Servers to Run Connectors Remotely*".

1.3. Overview of OpenICF Functionality

OpenICF provides many capabilities, including the following:

• Connector pooling

• Timeouts on all operations

• Search filtering

- Search and synchronization buffering and result streaming

- Scripting with Groovy, JavaScript, shell, and PowerShell

- Classloader isolation

- An independent logging API/SPI

- Java and .NET platform support

- Opt-in operations that support both simple and advanced implementations for the same CRUD operation

- A logging proxy that captures all API calls

- A Maven connector archetype to create connectors

Chapter 2
Using the OpenICF API

This chapter describes how to use the OpenICF API, which enables you to call OpenICF connector implementations from your application. The chapter demonstrates creating a connector facade, required for applications to access connectors, and then how to call the various OpenICF operations from your application.

2.1. Before You Start Using the OpenICF API

Before you can use an OpenICF connector in your application, you must download the OpenICF framework libraries, and the required connector bundles.

The easiest way to start using the OpenICF framework, from Java, is to use the sample Maven project file as a starting point. This sample project includes comprehensive comments about its use.

To use a .NET connector remotely, you must install the .NET remote connector server, as described in "Installing and Configuring a .NET Connector Server" in the *Integrator's Guide*. You must also download and install the specific connectors that you want to use from ForgeRock's BackStage site.

You can now start integrating the connector with your application.

2.2. About the Connector Facade

An application interacts with a connector through an instance of the `ConnectorFacade` class. The following diagram shows the creation and configuration of the connector facade. The components shown here are described in more detail in the sections that follow.
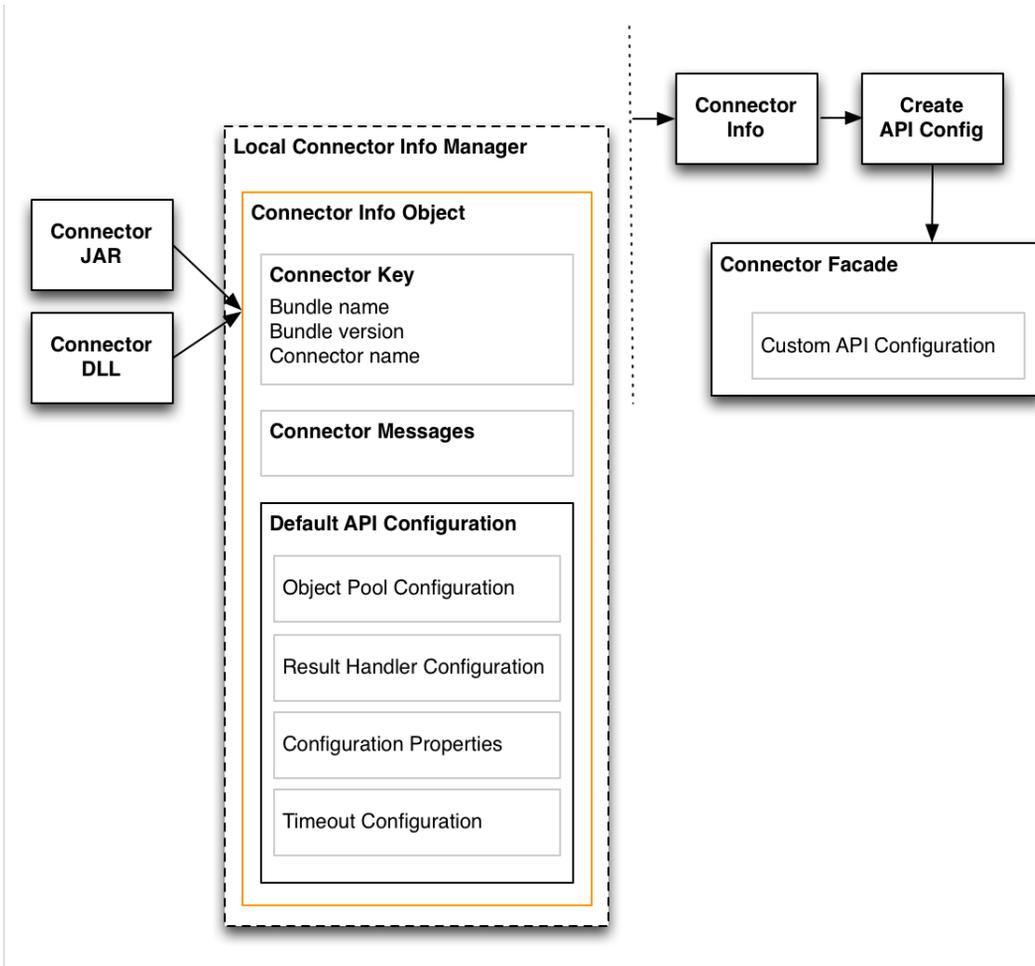
The connector facade is instantiated and configured in the following steps:

1. The application creates a `LocalConnectorInfoManager` instance (or instances) and adds the individual connector bundles (or assemblies).

   The `LocalConnectorInfoManager` processes these bundles or assemblies to instantiate a `ConnectorInfo` object.

   To be processed by the connector info manager, the connector bundle or assembly must have the following characteristics:

Java Connector Bundle

The `META-INF/MANIFEST.MF` file *must* include the following entries:

`ConnectorBundle-FrameworkVersion` - Minimum required OpenICF Framework version (either 1.1, 1.4, or 1.5)
`ConnectorBundle-Name` - Unique name of the connector bundle
`ConnectorBundle-Version` - Version of the connector bundle

The combination of the `ConnectorBundle-Name` and the `ConnectorBundle-Version` must be unique.

The connector bundle JAR must contain at least one class, that has the `ConnectorClass` annotation and implements the `Connector` interface.

.NET Connector Assembly

The `AssemblyInfo.cs` is used to determine the bundle version, from the `AssemblyVersion` property.

The bundle name is derived from the `Name` property of the assembly. For more information, see the corresponding Microsoft documentation.

> **Warning**
>
> If you change the name of your assembly, you must adjust the `bundleName` property in your connector configuration file, accordingly.

The connector assembly DLL must contain at least one class, that has the `ConnectorClassAttribute` attribute and implements the `Connector` interface.

2. For each connector, the `LocalConnectorInfoManager` processes the `MessageCatalog`, which contains the localized help and description messages for the configuration, and any log or error messages for the connector.

   Your application can use this information to provide additional help during the connector configuration process.

3. For each connector, the `LocalConnectorInfoManager` then processes the `ConfigurationClass`, to build the configuration properties for the connector.

4. Your application finds the connector info by its *connector key*. When the application has the connector info, it creates an API Configuration object that customises the following components:

   - Object pool configuration

   - Result handler configuration

   - Configuration properties

   - Timeout configuration

The API Configuration object is described in more detail in "The API Configuration Object".

5. The `ConnectorFacade` takes this customized API configuration object, determines which connector to use and how to configure it, and implements all of the OpenICF API operations.

2.3. The Connector Messages Object

The Connector Messages interface sets the message catalog for each connector, and enables messages to be localized. The interface has one method (`format()`), which formats a given message key in the current locale.

For more information, see the corresponding Javadoc.

2.4. The API Configuration Object

The API Configuration Object holds the runtime configuration of the connector facade instance. The OpenICF framework creates a default API Configuration Object inside the Connector Info Object. The application creates a copy of the API Configuration Object and customises it according to its requirements. The API Configuration Object includes the following components:

Object Pool Configuration

The object pool configuration specifies the pool configuration for poolable connectors only. Non-poolable connectors ignore this parameter. The object pool configuration includes the following parameters:

maxObjects

The maximum number of idle and active instances of the connector.

maxIdle

The maximum number of idle instances of the connector.

maxWait

The maximum time, in milliseconds, that the pool waits for an object before timing out. A value of `0` means that there is no timeout.

minEvictableIdleTimeMillis

The maximum time, in milliseconds, that an object can be idle before it is removed. A value of `0` means that there is no idle timeout.

minIdle

The minimum number of idle instances of the connector.

Results Handler Configuration

The results handler configuration defines how the OpenICF framework chains together the different results handlers to filter search results.

enableNormalizingResultsHandler

boolean

If the connector implements the attribute normalizer interface, you can enable this interface by setting this configuration property to `true`. If the connector does not implement the attribute normalizer interface, the value of this property has no effect.

enableFilteredResultsHandler

boolean

If the connector uses the filtering and search capabilities of the remote connected system, you can set this property to `false`. If the connector does not use the remote system's filtering and search capabilities (for example, the CSV file connector), you *must* set this property to `true`, otherwise the connector performs an additional, case-sensitive search, which can cause problems.

enableCaseInsensitiveFilter

boolean

By default, the filtered results handler (described previously) is case sensitive. If the filtered results handler is enabled this property allows you to enable case insensitive filtering. When case insensitive filtering is not enabled, a search will not return results unless the case matches exactly. For example, a search for `lastName = "Jensen"` will not match a stored user with `lastName : jensen`.

enableAttributesToGetSearchResultsHandler

boolean

By default, IDM determines which attributes that should be retrieved in a search. If the `enableAttributesToGetSearchResultsHandler` property is set to `true` the OpenICF framework removes all attributes from the READ/QUERY response, except for those that are specifically requested. For performance reasons, it is recommended that you set this property to `false` for local connectors, and to `true` for remote connectors.

Configuration Properties

The Configuration Properties object is built and populated by the framework as it parses the connectors configuration class.

Timeout Configuration

The timeout configuration enables you to configure timeout values per operation type. By default, there is no timeout configured for any operation type.

2.5. Creating the Connector Info Manager

You must initiate a specific connector info manager type, depending on whether your connector is local or remote. The following samples show how to create a local connector info manager and a remote connector info manager.

1. Create a `ConnectorInfoManager` and a `ConnectorKey` for the connector.

   The `ConnectorKey` uniquely identifies the connector instance. The `ConnectorKey` class takes a `bundleName` (the name of the Connector bundle), a `bundleVersion` (the version of the Connector bundle) and a connectorName (the name of the Connector)

   The `ConnectorInfoManager` retrieves a `ConnectorInfo` object for the connector by its connector key.

   Acquiring a Local Connector Info Object (Java)

   ```
   ConnectorInfoManagerFactory fact = ConnectorInfoManagerFactory.getInstance();
   File bundleDirectory = new File("/connectorDir/bundles/myconnector");
   URL url = IOUtil.makeURL(bundleDirectory,
       "/dist/org.identityconnectors.myconnector-1.0.jar");
   ConnectorInfoManager manager = fact.getLocalManager(url);
   ConnectorKey key = new ConnectorKey("org.identityconnectors.myconnector",
       "1.0", "MyConnector");
   ```

   Acquiring a Remote Connector Info Object (Java)

   ```
   ConnectorInfoManagerFactory fact = ConnectorInfoManagerFactory.getInstance();
   File bundleDirectory = new File("/connectorDir/bundles/myconnector");
   URL url = IOUtil.makeURL(bundleDirectory,
       "/dist/org.identityconnectors.myconnector-1.0.jar");
   ConnectorInfoManager manager = fact.getLocalManager(url);
   ConnectorKey key = new ConnectorKey("org.identityconnectors.myconnector",
       "1.0", "MyConnector");
   ```

2.6. Creating the Connector Facade

Applications access the connector API through a `ConnectorFacade` class, and interact with the connector through a `ConnectorFacade` instance.

The following steps describe how to create a `ConnectorFacade` in your application.

1. Create a `ConnectorInfoManager` and acquire the `ConnectorInfo` object for your connector, as described in the previous section.

2. From the `ConnectorInfo` object, create the default `APIConfiguration`.

```
APIConfiguration apiConfig = info.createDefaultAPIConfiguration();
```

3. Use the default `APIConfiguration` to set the `ObjectPoolConfiguration`, `ResultsHandlerConfiguration`, `ConfigurationProperties`, and `TimeoutConfiguration`.

```
ConfigurationProperties properties = apiConfig.getConfigurationProperties();
```

4. Set all of the `ConfigurationProperties` that you need for the connector, using `setPropertyValue().`

```
properties.setPropertyValue("host", SAMPLE_HOST);
properties.setPropertyValue("adminName", SAMPLE_ADMIN);
properties.setPropertyValue("adminPassword", SAMPLE_PASSWORD);
properties.setPropertyValue("useSSL", false);
```

5. Use the `newInstance()` method of the `ConnectorFacadeFactory` to create a new instance of the connector.

```
ConnectorFacade conn = ConnectorFacadeFactory.getInstance()
        .newInstance(apiConfig);
```

6. Validate that you have set up the connector configuration correctly.

```
conn.validate();
```

7. Use the new connector with the supported operations (described in the following sections).

```
conn.[authenticate|create|update|delete|search|...]
```

2.7. Checking the Schema and the Supported Operations

Different connectors support different subsets of the overall set of operations provided by OpenICF. When your connector is ready to use, you can use the `ConnectorFacade` to determine which operations your connector supports.

The quickest way to check whether an operation is supported is to determine whether that specific operation is part of the set of supported operations. The following sample test checks if the `CreateApiOp` is supported:

```
Set<Class< ? extends APIOperation>> ops = conn.getSupportedOperations();
return ops.contains(CreateApiOp.class);
```

Note that a connector might support a particular operation, only for specific object classes. For example, the connector might allow you to *create* a user, but not a group.

To be able to determine the list of supported operations for each object class, you need to check the schema. To determine whether the connector supports an operation for a specific object class, check the object class on which you plan to perform the operation, as shown in the following example.

```
Schema schema = conn.schema();
Set<ObjectClassInfo> objectClasses = schema.getObjectClassInfo();
Set<ObjectClassInfo> ocinfos = schema
        .getSupportedObjectClassesByOperation(CreateApiOp.class);

for(ObjectClassInfo oci : objectClasses) {
    // Check that the operation is supported for your object class.
    if (ocinfos.contains(ocinfo)) {
        // object class is supported
    }
}
```

In addition to determining the supported operations for an object class, your application can check which attributes are *required* and which attributes are *allowed* for a particular object class. The `ObjectClassInfo` class contains this information as a set of `AttributeInfo` objects.

The following example shows how to retrieve the attributes for an object class.

```
Schema schema = conn.schema();
Set<ObjectClassInfo> objectClasses = schema.getObjectClassInfo();
for(ObjectClassInfo oci : objectClasses) {
    Set<AttributeInfo> attributeInfos = oci.getAttributeInfo();
    String type = oci.getType();
    if(ObjectClass.ACCOUNT_NAME.equals(type)) {
        for(AttributeInfo info : attributeInfos) {
            System.out.println(info.toString());
        }
    }
}
```

Using the schema object, you can obtain the following information:

- Object classes and their attributes

- Operation options per operation

The following example shows how to retrieve the schema as a list of `ObjectClass` objects, from the `ObjectClassInfo` class.

```
ObjectClass objectClass = new ObjectClass(objectClassInfo.getType());
```

2.7.1. Operation Options

Operation options provide an extension point to an operation, enabling you to request additional information from the application, for each operation. The connector framework includes a number of predefined operation options for the most common use cases. For example, the option `OP_ATTRIBUTES_TO_GET` enables you to specify a list of attributes that should be returned by an operation.

When you write a connector, you must define the operation options that your connector supports in the schema, so that the application knows which operation options are supported.

For a list of the predefined operation options, see the corresponding Javadoc.

2.7.2. OpenICF Special Attributes

OpenICF includes a number of *special* attributes, that all begin and end with `__` (for example `__NAME__`, and `__UID__`). These special attributes are essentially functional aliases for specific attributes or object types. The purpose of the special attributes is to enable a connector developer to create a contract regarding how a property can be referenced, regardless of the application that is using the connector. In this way, the connector can map specific object information between an arbitrary application and the resource, without knowing how that information is referenced in the application.

The special attributes are used extensively in the generic LDAP connector, which can be used with ForgeRock Directory Services (DS), Active Directory, OpenLDAP, and other LDAP directories. Each of these directories might use a different attribute name to represent the same type of information. For example, Active Directory uses `unicodePassword` and DS uses `userPassword` to represent the same thing, a user's password. The LDAP connector uses the special OpenICF `__PASSWORD__` attribute to abstract that difference.

For a list of the special attributes, see the corresponding Javadoc.

2.8. How the OpenICF Framework Manages Connector Instances

The OpenICF framework supports multiple *connector types*, based on the implementation of the `connector` interface, and the `configuration` interface. These two interfaces determine the following:

- Whether the connector instance is obtained from a pool or whether a new instance is created *for each operation*

- Whether the connector configuration instance is retained, and reused for each operation, (stateful configuration) or a new configuration instance is created for each operation (stateless).

Connector developers determine what type of connector to implement, assessing the best match for the resource to which they are connecting. The interaction between the `connector` and `configuration` interface implementations is described in detail in "Deciding on the Connector Type". This section illustrates how the OpenICF framework manages connector instantiation, depending on the connector type.

2.8.1. Connector Instantiation for a Stateless, Non-Poolable Connector

The most basic connector has a stateless configuration, and is not pooled. A basic connector is initialized as follows:

1. The application calls an operation (for example, CREATE) on the connector facade.

2. The OpenICF framework creates a new *configuration instance*, and initializes it with its configuration properties.

3. When the framework has the configuration instance, with all the attributes in the configuration set, the framework creates a new *connector instance*, and initializes it, with the configuration that has been set.

4. The framework executes the operation (for example, CREATE) on the connector instance.

5. The connector instance executes the operation on the resource.

6. The framework calls the `dispose()` method to release all resources that the connector instance was using.

The following illustration shows the initialization process for a basic connector, and references the numbered steps in the preceding list.
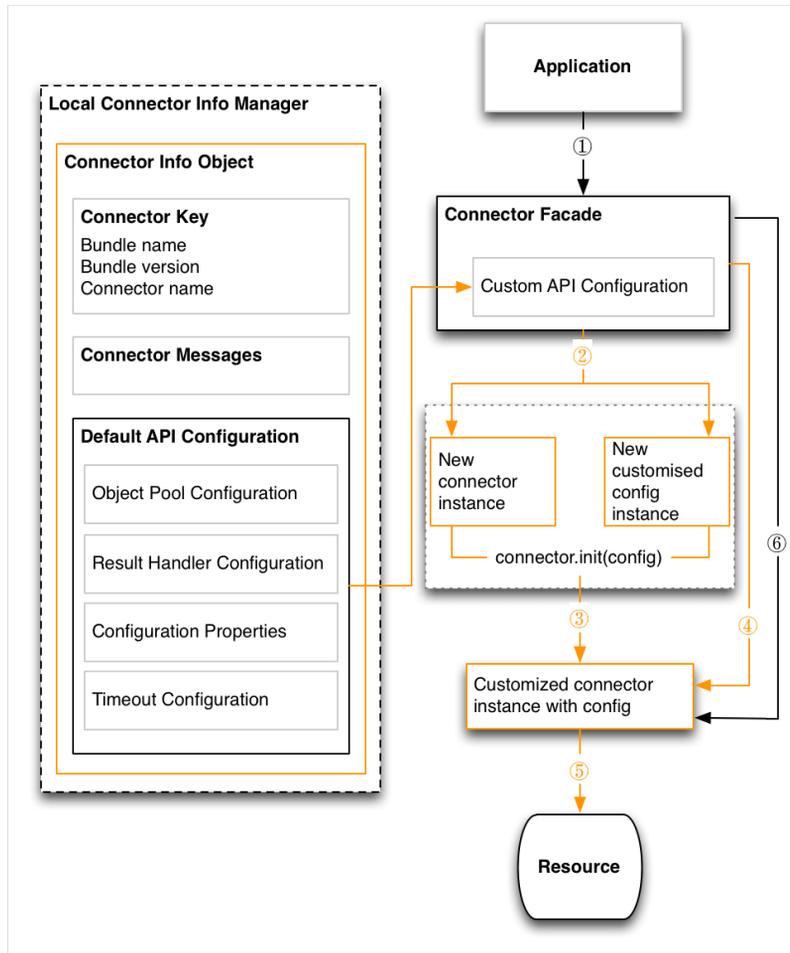
2.8.2. Connector Instantiation for a Stateless, Poolable Connector

The second connector type has a stateless configuration, but can be pooled. A stateless, poolable connector is instantiated as follows:

1. The application calls an operation (for example, CREATE) on the connector facade.

2. The OpenICF framework calls on the object pool, to borrow a *live* connector instance to execute the operation.

   If the object pool has an idle connector instance available, the framework *borrows* that one instance (step 5a in the illustration that follows).

The framework calls the `checkAlive` method on the customized connector instance with its configuration, to check if the instance that was borrowed from the pool is still alive, and ready to execute the operation. If the instance is no longer alive and ready, the framework disposes of the instance and borrows another one.

The thread that borrows the object has exclusive access to that connector instance, that is, it is thread-safe.

3. If the object pool has no idle connector instances, the pool creates a new connector instance.

4. The framework creates a new *configuration instance*, and initializes it with its configuration properties.

5. The framework initializes the borrowed connector instance, with the configuration that has been set.

6. The framework executes the operation (for example, CREATE) on the connector instance.

7. The connector instance executes the operation on the resource.

8. When the operation is complete, the framework releases the connector instance back into the pool. No `dispose()` method is called.

The following illustration shows the initialization process for a stateless, poolable connector, and references the numbered steps in the preceding list.
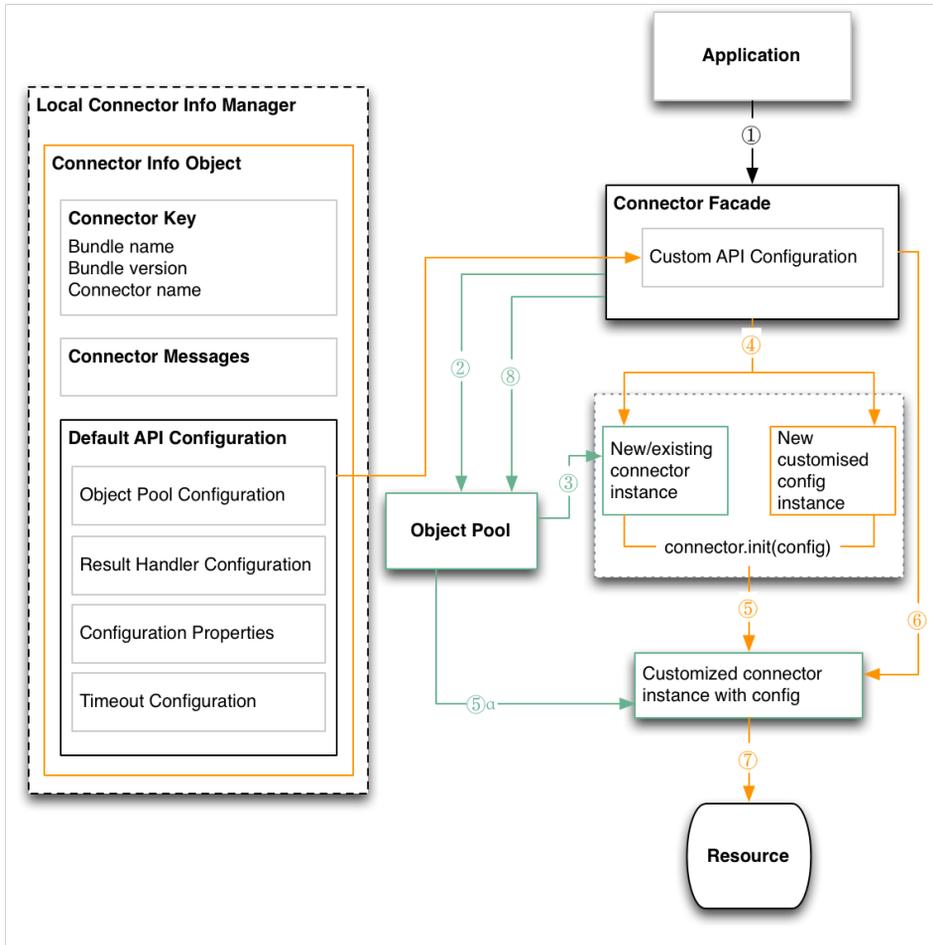
2.8.3. Connector Instantiation for a Stateful, Non-Poolable Connector

The third connector type has a stateful configuration, and cannot be pooled. A stateful, non-poolable connector is instantiated as follows:

1. The OpenICF framework creates a new *configuration instance*, initializes it with its configuration properties, and stores it in the connector facade, before any operations are called.

   This single configuration instance is shared between multiple threads. The framework does not guarantee isolation, so connector developers must ensure that their implementation is thread-safe.

2. The application calls an operation (for example, CREATE) on the connector facade.

3. The OpenICF framework creates a new connector instance, and calls the `init()` method on that connector instance, with the stored configuration. The framework initializes the connector with the single configuration instance stored within the connector facade.

4. The framework executes the operation (for example, CREATE) on the connector instance.

5. The connector instance executes the operation on the resource.

6. The framework calls the `dispose()` method to release all resources that the connector instance was using.

   Note that the customized config instance remains in the connector facade, and is reused for the next operation.

The following illustration shows the initialization process for a non-poolable connector, with a stateful configuration. The illustration references the numbered steps in the preceding list.
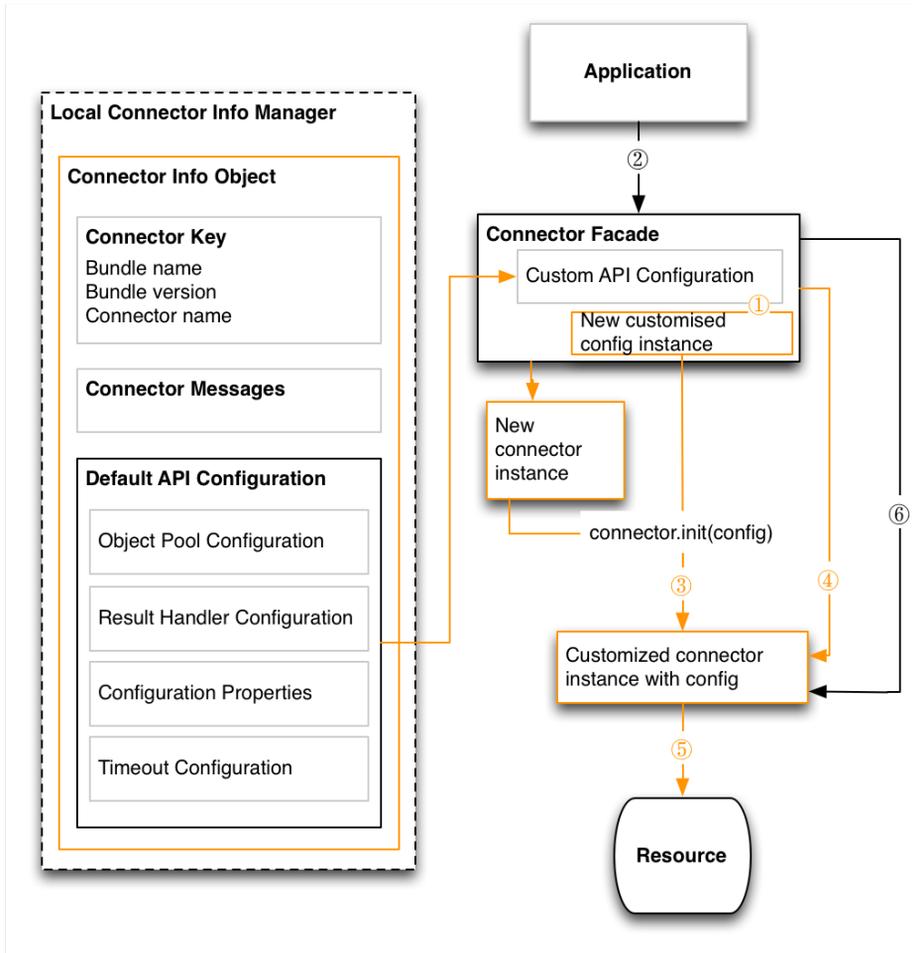
2.8.4. Connector Instantiation for a Stateful, Poolable Connector

The fourth connector type has a stateful configuration, and can be pooled. A stateful, poolable connector is instantiated as follows:

1. The OpenICF framework creates a new *configuration instance*, initializes it with its configuration properties, and stores it in the connector facade, before any operations are called.

    This single configuration instance is shared between multiple threads. The framework does not guarantee isolation, so connector developers must ensure that their implementation is thread-safe.

2. The application calls an operation (for example, CREATE) on the connector facade.

3. The framework calls on the object pool, to borrow a *live* connector instance to execute the operation.

   If the object pool has an idle connector instance available, the framework *borrows* that one instance (step 5a in the illustration that follows).

   The framework calls the `checkAlive` method on the customized connector instance with its configuration, to check if the instance that was borrowed from the pool is still alive, and ready to execute the operation. If the instance is no longer alive and ready, the framework disposes of the instance and borrows another one.

   The thread that borrows the object has exclusive access to that connector instance, that is, it is thread-safe.

4. If the object pool has no idle connector instances, the pool creates a new connector instance.

5. The framework initializes the borrowed connector instance, with the stored configuration.

6. The framework executes the operation (for example, CREATE) on the connector instance.

7. The connector instance executes the operation on the resource.

8. When the operation is complete, the framework releases the connector instance back into the pool. No `dispose()` method is called.

The following illustration shows the initialization process for a stateful, poolable connector, and references the numbered steps in the preceding list.
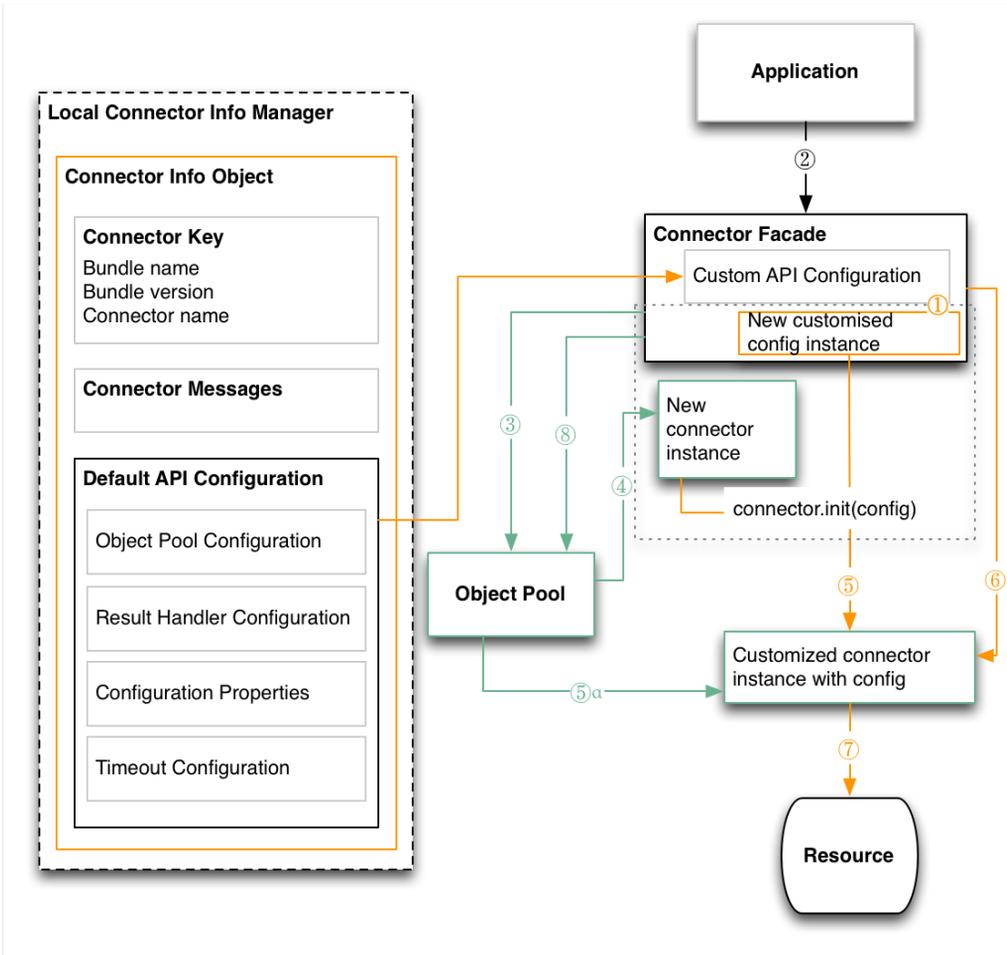
Chapter 3

Using OpenICF Connector Servers to Run Connectors Remotely

A *Connector Server* enables your application to run provisioning operations on a connector bundle that is deployed on a remote system. A key feature of connector servers is the ability to run connector bundles that are written in C# on a .NET platform, and to access them over the network from a Java application.

Connector servers are available for both the Java and .NET platforms.

- A .NET connector server allows Java applications to access C# connector bundles. You deploy the C# connector bundles on the .NET connector server. Your Java application can then communicate with the .NET connector server over the network. The .NET connector server acts as a proxy to provide authenticated access for Java applications to the C# connector bundles that are deployed within the connector server.

- A Java connector server allows you to execute a Java connector bundle in a different JVM from your application. You can also run a Java connector on a different host for performance reasons. You might choose to use a Java connector server with a Java remote connector server to avoid the possibility of crashing an application JVM, due to a fault in a JNI-based connector.

For information on installing and configuring a .NET connector server, see "Installing and Configuring a .NET Connector Server" in the *Integrator's Guide*. For information on installing and configuring a Java connector server, see "Installing and Configuring a Remote Java Connector Server" in the *Integrator's Guide*.

3.1. Accessing Connector Servers Over SSL

In production environments, your application should access either the .NET or remote Java Connector Server over SSL.

To configure secure access to the connector server, follow these steps:

1. Configure the Connector Server to use an SSL certificate.

   Edit the `ConnectorServer.properties` file (in the `/path/to/openicf/conf` folder), to set `connectorserver.usessl=true`.

2. Configure the Connector Server to provide SSL sockets.

If you use self-signed or other certificates that Java cannot verify automatically, either import these certificates into the `$JAVA_HOME/lib/security/cacerts` directory, or use Java properties to specify your own, properly configured, truststore:

- `-Djavax.net.ssl.trustStore=/path/to/myApp_cacerts`

- `-Djavax.net.ssl.trustStorePassword=changeit`

Chapter 4
Implementing the OpenICF SPI

This chapter describes the OpenICF SPI, which enables you to create connectors that are compatible with the OpenICF framework.

The SPI includes a number of interfaces, but you need only implement those that are supported by the target resource to which you are connecting. For information about how to get started with writing connectors, see "*Writing Java Connectors*" and "*Writing Scripted Connectors With the Groovy Connector Toolkit*".

The order in which you implement your connector is as follows:

1. Decide on the connector type (see "Deciding on the Connector Type".

2. Implement the configuration interface (see "Implementing the Configuration Interface".

3. Implement the connector interface (see "Implementing the Connector Interface".

4. Implement the operation interfaces (see "Implementing the Operation Interfaces".

4.1. Deciding on the Connector Type

OpenICF supports multiple connector types, based on the implementation of the `connector` interface, and the `configuration` interface. These two interfaces determine whether the connector can be pooled, and whether its configuration is stateful. Before you begin developing your connector, decide on the *connector type*, based on the system to which you are connecting. For an overview of how the OpenICF framework manages each connector type, see "How the OpenICF Framework Manages Connector Instances".

This section outlines the different connector types.

Connector

> The basic connector is a *non-poolable* connector. Each operation is executed on a new instance of the connector. OpenICF creates a new instance of the Connector class and uses a new or existing instance of the connector configuration to initialise the instance before the operation is executed. After the execution, OpenICF disposes of the connector instance.

Poolable Connector

>   Before an operation is executed, an existing connector instance is pulled from the Connector Pool. If there is no existing instance, a new instance is created. After the operation execution, the Connector instance is released and placed back into pool.

>   The OpenICF framework pools *instances* of a poolable connector, rather than pooling connections within the connector.

Configuration

>   For a basic (non-stateful) configuration, each time the configuration is used (when an operation is validated or a new connector instance is initialised, a new Configuration instance is created and configured with the Configuration properties.

Stateful Configuration

>   With a stateful configuration, the configuration instance is created only once and is used until the Facade or Connector Pool that is associated with the Configuration is disposed of.

The following table illustrates how these elements combine to determine the connector type.

Connector Types

	Connector	Poolable Connector
Configuration	Entirely stateless combination. A new Configuration and Connector instance are created for each operation.	Connector initialisation is an expensive operation, so it is preferable to keep connector instances in a pool. A new configuration is required only when a new connector instance is added to the pool.
Stateful Configuration	The configuration can be used to make the heavy resource initialisation. The less intensive connector instance can then execute the operation.	The configuration must be shared between the instances in the same pool and the connector initialisation is expensive.

For detailed information on how the OpenICF framework manages each of these connector types, see "How the OpenICF Framework Manages Connector Instances".

4.2. Implementing the Configuration Interface

The OpenICF connector framework uses the configuration interface implementation to build the *configuration properties* inside the API configuration.

The configuration interface implementation includes the required information to enable the connector to connect to the target system, and to perform its operations. The configuration interface

implements getters and setters for each of its defined properties. It also provides a validate method that your application can use to check whether all the required properties are available, and valid, before passing them to the connector.

The configuration interface has three methods:

- `setConnectorMessages(ConnectorMessages messages)` sets the message catalog instance, that enables the connector to provide localized messages.

  The message catalog is defined in the file `Messages.properties`, and can be localized as required by appending the locale to the file name, for example, `Messages_fr.properties`.

  For more information on the message catalog, see "The Connector Messages Object".

- `getConnectorMessages()` returns the message catalog that is set by setConnectorMessages(ConnectorMessages)

- `validate()` checks that all the required properties have been set and that their values are valid

  The purpose of this method is to test that the configuration that the application provides to your connector is valid.

Each property that is declared is not necessarily required. If a property is required, it must be included in the `ConfigurationProperty` annotation.

The `ConfigurationProperty` annotation (Java) or attribute (.NET) enables you to add custom meta information to properties. The OpenICF framework scans the meta information and collects this information to build the `ConfigurationProperties` object inside the `APIConfiguration`. The following meta information can be provided:

Element	Description	Implementation in Java	Implementation in C#
order	The order in which this property is displayed		
helpMessageKey	Enables you to change the default help message key	*propertyName*.help	*help_propertyName*
displayMessageKey	Enables you to change the default display message key	*propertyName.display*	*display_propertyName*
groupMessageKey	Enables you to change the default group message key	*propertyName.group*	*group_propertyName*
confidential	Indicates that this is a confidential property and that its value should be encrypted by the application when persisted		

Element	Description	Implementation in Java	Implementation in C#
required	Boolean, indicates whether the property is required		
operations	The array of operations that require this property		

The following examples show how the meta information is provided, in both Java and C#.

Stateless Configuration Implementation (Java)

```java
public class SampleConfiguration extends AbstractConfiguration  {

    /**
     * {@inheritDoc}
     */
    public void validate() {
    }

    @ConfigurationProperty(
        order = 1,
        helpMessageKey = "passwordFieldName.help",
        displayMessageKey = "passwordFieldName.display",
        groupMessageKey = "authenticateOp.group",
        confidential = false,
        required = false,
        operations = {AuthenticateOp.class,CreateOp.class}
    )
    public String getPasswordFieldName() {
        return passwordFieldName;
    }

    public void setPasswordFieldName(String value) {
        passwordFieldName = value;
    }
}
```

Stateful Configuration Implementation (Java)

```java
public class SampleConfiguration extends AbstractConfiguration
    implements StatefulConfiguration {

    /**
     * {@inheritDoc}
     */
    public void release() {
    }

    /**
     * {@inheritDoc}
     */
    public void validate() {
    }
}
```

Stateless Configuration Implementation (C#)

```csharp
public class ActiveDirectoryConfiguration : AbstractConfiguration
    {

        [ConfigurationProperty(
            Order = 1,
            HelpMessageKey = "help_PasswordFieldName",
            DisplayMessageKey = "display_PasswordFieldName",
            GroupMessageKey = "group_PasswordFieldName",
            Confidential = false,
            Required = false,
            OperationTypes = new[] { typeof(AuthenticateOp) })
        ]
        public String PasswordFieldName
        { get; set; }

        public override void Validate()
        {
            throw new NotImplementedException();
        }
    }
```

Stateful Configuration Implementation (C#)

```csharp
public class ActiveDirectoryConfiguration : AbstractConfiguration,
    StatefulConfiguration
    {

        public override void Validate()
        {
            throw new NotImplementedException();
        }

        public void Release()
        {
            throw new NotImplementedException();
        }
    }
```

4.2.1. Validate Operation

The validate operation validates the connector configuration. A valid configuration is one that is *ready to be used* by the connector.

A configuration that is *ready*, has the following characteristics:

- It is complete, that is all required properties are present and have values

- All property values are well-formed, that is, they are in the expected range and have the expected format

4.2.1.1. ValidateApiOp

The validate operation returns a `ConfigurationException` in the following situations:

- The Framework version is not compatible with the connector

- The connector does not have the required attributes in `MANIFEST.MF`

- The `ConfigurationProperties` cannot be merged into the configuration

Implementation of the valid operation, at the API Level

```java
@Test
    public void ValidateTest() {
    logger.info("Running Validate Test");
    final ConnectorFacade facade = createConnectorFacade(BasicConnector.class, null);
    facade.validate();
    }
```

4.2.2. Validate SPI Implementation

The `validate()` method of the configuration operation must return one of the following:

- `RuntimeException` if the configuration is not valid

- `NullPointerException` if a required configuration property is null

- `IllegalArgumentException` if a required configuration property is blank

Implementation of the validate method

```
public void validate() {
    if (StringUtil.isBlank(host)) {
        throw new IllegalArgumentException("Host User cannot be null or empty.");
    }

    Assertions.blankCheck(remoteUser, "remoteUser");

    Assertions.nullCheck(password, "password");
}
```

4.2.3. Supported Configuration Types

The OpenICF framework supports a limited number of configuration property types. This limitation is necessary, because OpenICF must serialise and deserialize the configuration property values when sending them over the network.

You can use any of the following types, or an array of these types. Lists of types are not supported.

```
String.class
long.class
Long.class
char.class
Character.class
double.class
Double.class
float.class
Float.class
int.class
Integer.class
boolean.class
Boolean.class
URI.class
File.class
GuardedByteArray.class
GuardedString.class
Script.class
```

```
    typeof(string),
    typeof(long),
    typeof(long?),
    typeof(char),
    typeof(char?),
    typeof(double),
    typeof(double?),
    typeof(float),
    typeof(float?),
    typeof(int),
    typeof(int?),
    typeof(bool),
    typeof(bool?),
    typeof(Uri),
    typeof(FileName),
    typeof(GuardedByteArray),
    typeof(GuardedString),
    typeof(Script)
```

The framework introspects the implemented configuration class and adds all properties that have a `set/get` method to the `ConfigurationProperties` object.

The `ConfigurationClass` annotation (Java) or attribute (.NET) provides additional information to the OpenICF framework about the configuration class. The following information is provided:

Element	Description
privateProperty	If this is set, the property is hidden from the application, and the application cannot set the property through the `APIConfiguration`.
skipUnsupported	If the type of an added property is not supported, the framework throws an exception. To avoid the exception, set the value of `skipUnsupported` to `true`.

ConfigurationClass Annotation in Java

```
@ConfigurationClass(ignore = { "privateProperty", "internalProperty" }, skipUnsupported = true)
```

ConfigurationClass Attribute in C#

```
[ConfigurationClass(Ignore = { "privateProperty", "internalProperty" }, SkipUnsupported = true)]
```

4.3. Implementing the Connector Interface

The connector interface declares a connector, and manages its life cycle. You *must* implement the `connector` interface. A typical connector lifecycle is as follows:

- The connector creates a connection to the target system.

- Any operations implemented in the connector are called.

- The connector discards the connection and disposes of any resources it has used.

The `connector` interface has only three methods:

- `init(Configuration)` initializes the connector with its configuration

- `getConfiguration()` returns the configuration that was passed to `init(Configuration)`

- `dispose()` disposes of any resources that the connector uses.

The `ConnectorClass`, which is the implementation of the connector interface, must have the `ConnectorClass` annotation (Java) or attribute (.NET) so that the OpenICF framework can find the connector class. The following table shows the elements within the connector class.

Element	Description		
configurationClass	The configuration class for the connector.		
displayNameKey	A key in the message catalog that holds a human readable name for the connector.		
categoryKey	The category to which the connector belongs, such as LDAP, or DB.		
messageCatalogPaths	The resource path(s) to the message catalog. If multiple paths are provided, the message catalogs are collated. By default, if no path is specified, the `connector-package.Messages.properties` is used		

The following examples show the connector interface implementation, in Java and C#.

Connector Interface Implementation in Java

```
@ConnectorClass(
    displayNameKey = "Sample.connector.display",
    configurationClass = SampleConfiguration.class)
public class SampleConnector implements Connector...
```

Connector Interface Implementation in C#

```
[ConnectorClass(
    "connector_displayName",
    typeof (SampleConfiguration)
    ]
public class SampleConnector : Connector ...
```

4.3.1. Implementing a Poolable Connector Interface

Certain connectors support the ability to be pooled. For a pooled connector, OpenICF maintains a pool of connector instances and reuses these instances for multiple provisioning and reconciliation operations. When an operation must be executed, an existing connector instance is taken from the connector pool. If no connector instance exists, a new instance is initialized. When the operation has been executed, the connector instance is released back into the connector pool, ready to be used for a subsequent operation.

For an unpooled connector, a new connector instance is initialized for every operation. When the operation has been executed, OpenICF disposes of the connector instance. Because the initialization of a connector is an expensive operation, reducing the number of connector initializations can substantially improve performance.

The following connection pooling configuration parameters can be set:

maxObjects

> The maximum number of connector instances in the pool (both idle and active). The default value is 10 instances.

maxIdle

> The maximum number of idle connector instances in the pool. The default value is 10 idle instances.

maxWait

> The maximum period to wait for a free connector instance to become available before failing. The default period is 150000 milliseconds, or 15 seconds.

minEvictableIdleTimeMillis

> The minimum period to wait before evicting an idle connector instance from the pool. The default period is 120000 milliseconds, or 12 seconds.

minIdle

> The minimum number of idle connector instances in the pool. The default value is 1 instance.

A `PoolableConnector` extends the connector interface with the `checkAlive()` method. You should use a `PoolableConnector` when the `init(Configuration)` method is so expensive that it is worth keeping the connector instance in a pool and reusing it between operations. When an existing connector instance is pooled, the framework calls the `checkAlive()` method. If this method throws an error, the framework discards it from the pool and obtains another instance, or creates a new connector instance and calls the `init()` method. The `checkAlive()` method is used to make sure that the instance in the pool is still operational.

4.4. Implementing the Operation Interfaces

The SPI provides several operations. The subset of operations that you implement will depend on the target resource to which you are connecting. Each operation interface defines an action that the connector can perform on the target resource.

The following sections describe the operation interfaces that are provided by the SPI, and provide examples of how they can be implemented in your connector. The sections include the API- and SPI-level rules for each operation.

4.4.1. Authenticate Operation

The authenticate operation authenticates an object on the target system, based on two parameters, usually a unique identifier (username) and a password. If possible, your connector should try to authenticate these credentials natively.

If authentication fails, the connector should throw a runtime exception. The exception must be an `IllegalArgumentException` or, if a native exception is available and is of type `RuntimeException`, that native runtime exception. If the native exception is not a `RuntimeException`, it should be wrapped in a `RuntimeException`, and then thrown.

The exception should provide as much detail as possible for logging problems and failed authentication attempts. Several exceptions are provided in the `exceptions` package, for this purpose. For example, one of the most common authentication exceptions is the `InvalidPasswordException`.

For more information about the common exceptions provided in the OpenICF framework, see "Common Exceptions".

4.4.1.1. Using the OpenICF Authenticate Operation

This section shows how your application can use the framework's `authentication` operation, and how to write a unit test for this operation, when you are developing your connector.

The `authentication` operation throws a `RuntimeException` if the credentials do not pass authentication, otherwise returns the `UID`.

Sample Unit Test for the Authentication Operation (Java)

```
@Test
public void authenticateTest() {
    logger.info("Running Authentication Test");
    final ConnectorFacade facade = createConnectorFacade(BasicConnector.class, null);
    final OperationOptionsBuilder builder = new OperationOptionsBuilder();
    Uid uid =
        facade.authenticate(ObjectClass.ACCOUNT, "username", new GuardedString("Passw0rd"
        .toCharArray()), builder.build());
    Assert.assertEquals(uid.getUidValue(), "username");
}
```

4.4.1.2. Implementing the Authenticate Operation in Your Connector

To implement the `authenticate` operation in your connector, add the `AuthenticateOp` interface to your connector class, for example:

```
@ConnectorClass(
    displayNameKey = "Sample.connector.display",
    configurationClass = SampleConfiguration.class)
public class SampleConnector implements Connector, AuthenticateOp...
```

For more information, see the AuthenticateOp JavaDoc.

The SPI provides the following detailed exceptions:

- UnknownUidException - the UID does not exist on the resource

  `(org.identityconnectors.framework.common.exceptions.UnknownUidException)`

- ConnectorSecurityException - base exception for all security-related exceptions

  `(org.identityconnectors.framework.common.exceptions.ConnectorSecurityException)`

- InvalidCredentialException - generic invalid credential exception that should be used if the specific error cannot be obtained

  `(org.identityconnectors.framework.common.exceptions.UnknownUidException)`

- InvalidPasswordException - the password provided is incorrect

  `(org.identityconnectors.framework.common.exceptions.InvalidPasswordException)`

- PasswordExpiredException - the password is correct, but has expired

  `(org.identityconnectors.framework.common.exceptions.PasswordExpiredException)`

- PermissionDeniedException - the user can be identified but does not have permission to authenticate

```
(org.identityconnectors.framework.common.exceptions.PermissionDeniedException)
```

Implementation of the Authentication Operation, at the SPI Level

```java
public Uid authenticate(final ObjectClass objectClass, final String userName,
        final GuardedString password, final OperationOptions options) {
    if (ObjectClass.ACCOUNT.equals(objectClass)) {
        return new Uid(userName);
    } else {
        logger.warn("Authenticate of type {0} is not supported", configuration
                .getConnectorMessages().format(objectClass.getDisplayNameKey(),
                        objectClass.getObjectClassValue())));
        throw new UnsupportedOperationException("Authenticate of type"
                + objectClass.getObjectClassValue() + " is not supported");
    }
}
```

4.4.2. Create Operation

The create operation interface enables the connector to create objects on the target system. The operation includes one method (`create()`). The method takes an `ObjectClass`, and any provided attributes, and creates the object and its UID. The connector must return the UID so that the caller can refer to the created object.

The connector should make a best effort to create the object, and should throw an informative `RuntimeException`, indicating to the caller why the operation could not be completed. Defaults can be used for any required attributes, as long as the defaults are documented.

The UID is never passed in with the attribute set for this method. If the resource supports a mutable UID, you can create a resource-specific attribute for the ID, such as `unix_uid`.

If the `create` operation is only partially successful, the connector should attempt to roll back the partial change. If the target system does not allow this, the connector should report the partial success of the create operation and throw a `RetryableException`. For example:

```java
public static RetryableException wrap(final String message, final Uid uid) {
    return new RetryableException(message, new AlreadyExistsException().initUid(Assertions
    .nullChecked(uid, "Uid")));
}
```

4.4.2.1. Using the OpenICF Create Operation

The following exceptions are thrown by the Create API operation:

- `IllegalArgumentException` - if `ObjectClass}` is missing, or if elements of the set produce duplicate values of `Attribute#getName()`

- `NullPointerException` - if the `createAttributes` parameter is `null`

- `RuntimeException` - if the `Connector` SPI throws a native exception

Consumption of the Create Operation, at the API Level

```
@Test
public void createTest() {
    logger.info("Running Create Test");
    final ConnectorFacade facade = createConnectorFacade(BasicConnector.class, null);
    final OperationOptionsBuilder builder = new OperationOptionsBuilder();
    Set<Attribute> createAttributes = new HashSet<Attribute>();
    createAttributes.add(new Name("Foo"));
    createAttributes.add(AttributeBuilder.buildPassword("Password".toCharArray()));
    createAttributes.add(AttributeBuilder.buildEnabled(true));
    Uid uid = facade.create(ObjectClass.ACCOUNT, createAttributes, builder.build());
    Assert.assertEquals(uid.getUidValue(), "foo");
}
```

4.4.2.2. Implementing the Create Operation in Your Connector

The SPI provides the following detailed exceptions:

- `UnsupportedOperationException` - the create operation is not supported for the specified object class

- `InvalidAttributeValueException` - a required attribute is missing, an attribute is present that cannot be created, or a provided attribute has an invalid value

- `AlreadyExistsException` - an object with the specified `Name` already exits on the target system

- `PermissionDeniedException` - the target resource will not allow the connector to perform the specified operation

- `ConnectorIOException, ConnectionBrokenException, ConnectionFailedException` - a problem as occurred with the connection

- `RuntimeException` - thrown if anything else goes wrong. You should try to throw a native exception in this case.

Implementation of the Create Operation, at the SPI Level

```java
public Uid create(final ObjectClass objectClass, final Set<Attribute> createAttributes,
        final OperationOptions options) {
    if (ObjectClass.ACCOUNT.equals(objectClass) || ObjectClass.GROUP.equals(objectClass)) {
        Name name = AttributeUtil.getNameFromAttributes(createAttributes);
        if (name != null) {
            // do real create here
            return new Uid(AttributeUtil.getStringValue(name).toLowerCase());
        } else {
            throw new InvalidAttributeValueException("Name attribute is required");
        }
    } else {
        logger.warn("Delete of type {0} is not supported", configuration.getConnectorMessages()
                .format(objectClass.getDisplayNameKey(), objectClass.getObjectClassValue()));
        throw new UnsupportedOperationException("Delete of type"
                + objectClass.getObjectClassValue() + " is not supported");
    }
}
```

4.4.3. Delete Operation

The delete operation interface enables the connector to delete an object on the target system. The operation includes one method (`delete()`). The method takes an `ObjectClass`, a `Uid`, and any operation options.

The connector should call the native delete methods to remove the object, specified by its unique ID.

4.4.3.1. Using the OpenICF Delete Operation

The following exceptions are thrown by the Delete API operation:

• `UnknownUidException` - the UID does not exist on the resource

Consumption of the Delete Operation, at the API Level

```java
@Test
public void deleteTest() {
    logger.info("Running Delete Test");
    final ConnectorFacade facade = createConnectorFacade(BasicConnector.class, null);
    final OperationOptionsBuilder builder = new OperationOptionsBuilder();
    facade.delete(ObjectClass.ACCOUNT, new Uid("username"), builder.build());
}
```

4.4.3.2. Implementing the Delete Operation in Your Connector

Implementation of the Delete Operation, at the SPI Level

```
public void delete(final ObjectClass objectClass, final Uid uid, final OperationOptions options) {
    if (ObjectClass.ACCOUNT.equals(objectClass) || ObjectClass.GROUP.equals(objectClass)) {
        // do real delete here
    } else {
        logger.warn("Delete of type {0} is not supported", configuration.getConnectorMessages()
                .format(objectClass.getDisplayNameKey(), objectClass.getObjectClassValue()));
        throw new UnsupportedOperationException("Delete of type"
                + objectClass.getObjectClassValue() + " is not supported");
    }
}
```

4.4.4. Resolve Username Operation

The resolve username operation enables the connector to resolve an object to its UID, based on its username. This operation is similar to the simple authentication operation. However, the resolve username operation does not include a password parameter, and does not attempt to authenticate the credentials. Instead, it returns the UID that corresponds to the supplied username.

The implementation must, however, validate the username (that is, the connector must throw an exception if the username does not correspond to an existing object). If the username validation fails, the the connector should throw a runtime exception, either an `IllegalArgumentException` or, if a native exception is available and is of type `RuntimeException`, simply throw that exception. If the native exception is not a `RuntimeException`, it should be wrapped in a `RuntimeException`, and then thrown.

The exception should provide as much detail as possible for logging problems and failed attempts. Several exceptions are provided in the `exceptions` package, for this purpose. For example, one of the most common exceptions is the `UnknownUidException`.

4.4.4.1. Using the OpenICF Resolve Username Operation

The operation throws a `RuntimeException` if the username validation fails, otherwise returns the `UID`.

Consumption of the ResolveUsername operation, at the API Level

```
@Test
public void resolveUsernameTest() {
    logger.info("Running ResolveUsername Test");
    final ConnectorFacade facade = createConnectorFacade(BasicConnector.class, null);
    final OperationOptionsBuilder builder = new OperationOptionsBuilder();
    Uid uid = facade.resolveUsername(ObjectClass.ACCOUNT, "username", builder.build());
    Assert.assertEquals(uid.getUidValue(), "username");
}
```

4.4.4.2. Implementing the Resolve Username Operation in Your Connector

The SPI provides the following detailed exceptions:

• UnknownUidException - the UID does not exist on the resource

Implementation of the ResolveUsername Operation, at the SPI Level

```
public Uid resolveUsername(final ObjectClass objectClass, final String userName,
        final OperationOptions options) {
    if (ObjectClass.ACCOUNT.equals(objectClass)) {
        return new Uid(userName);
    } else {
        logger.warn("ResolveUsername of type {0} is not supported", configuration
                .getConnectorMessages().format(objectClass.getDisplayNameKey(),
                        objectClass.getObjectClassValue()));
        throw new UnsupportedOperationException("ResolveUsername of type"
                + objectClass.getObjectClassValue() + " is not supported");
    }
}
```

4.4.5. Schema Operation

The Schema Operation interface enables the connector to describe the types of objects that it can handle on the target system, and the operations and options that the connector supports for each object type.

The operation has one method, `schema()`, which returns the types of objects on the target system that the connector supports. The method should return the object class name, its description, and a set of attribute definitions.

The implementation of this operation includes a mapping between the native object class and the corresponding connector object. The special `Uid` attribute should not be returned, because it is not a true attribute of the object, but a reference to it. For more information about special attributes in OpenICF, see "OpenICF Special Attributes".

If your resource object class has a writable unique ID attribute that is different to its `Name`, your schema should contain a resource-specific attribute that represents this unique ID. For example, a Unix account object might contain a `unix_uid`.

4.4.5.1. Using the OpenICF Schema Operation

Consumption of the Schema Operation, at the API Level

```
@Test
public void schemaTest() {
    logger.info("Running Schema Test");
    final ConnectorFacade facade = createConnectorFacade(BasicConnector.class, null);
    Schema schema = facade.schema();
    Assert.assertNotNull(schema.findObjectClassInfo(ObjectClass.ACCOUNT_NAME));
}
```

4.4.5.2. Implementing the Schema Operation in Your Connector

Implementation of the SchemaOp operation, at the SPI Level

```
public Schema schema() {
    if (null == schema) {
        final SchemaBuilder builder = new SchemaBuilder(BasicConnector.class);
        // Account
        ObjectClassInfoBuilder accountInfoBuilder = new ObjectClassInfoBuilder();
        accountInfoBuilder.addAttributeInfo(Name.INFO);
        accountInfoBuilder.addAttributeInfo(OperationalAttributeInfos.PASSWORD);
        accountInfoBuilder.addAttributeInfo(PredefinedAttributeInfos.GROUPS);
        accountInfoBuilder.addAttributeInfo(AttributeInfoBuilder.build("firstName"));
        accountInfoBuilder.addAttributeInfo(AttributeInfoBuilder.define("lastName")
                .setRequired(true).build());
        builder.defineObjectClass(accountInfoBuilder.build());

        // Group
        ObjectClassInfoBuilder groupInfoBuilder = new ObjectClassInfoBuilder();
        groupInfoBuilder.setType(ObjectClass.GROUP_NAME);
        groupInfoBuilder.addAttributeInfo(Name.INFO);
        groupInfoBuilder.addAttributeInfo(PredefinedAttributeInfos.DESCRIPTION);
        groupInfoBuilder.addAttributeInfo(AttributeInfoBuilder.define("members").setCreatable(
                false).setUpdateable(false).setMultiValued(true).build());

        // Only the CRUD operations
        builder.defineObjectClass(groupInfoBuilder.build(), CreateOp.class, SearchOp.class,
                UpdateOp.class, DeleteOp.class);

        // Operation Options
        builder.defineOperationOption(OperationOptionInfoBuilder.buildAttributesToGet(),
                SearchOp.class);

        // Support paged Search
        builder.defineOperationOption(OperationOptionInfoBuilder.buildPageSize(),
                SearchOp.class);
        builder.defineOperationOption(OperationOptionInfoBuilder.buildPagedResultsCookie(),
                SearchOp.class);

        // Support to execute operation with provided credentials
        builder.defineOperationOption(OperationOptionInfoBuilder.buildRunWithUser());
        builder.defineOperationOption(OperationOptionInfoBuilder.buildRunWithPassword());
```

```
        schema = builder.build();
    }
    return schema;
}
```

4.4.6. Script On Connector Operation

The script on connector operation runs a script in the environment of the connector. This is different to the script on resource operation, which runs a script on the target resource that the connector manages.

The corresponding API operation (`scriptOnConnectorApiOp`) provides a minimum contract to which the connector must adhere. (See the javadoc for more information). If you do not implement the `scriptOnConnector` interface in your connector, the framework provides a default implementation. If you intend your connector to provide more to the script than what is required by this minimum contract, you must implement the `scriptOnConnectorOp` interface.

4.4.6.1. Using the OpenICF Script on Connector Operation

The API operation allows an application to run a script in the context of any connector.

This operation runs the script in the same JVM or .Net Runtime as the connector. That is, if you are using a local framework, the script runs in your JVM. If you are connected to a remote framework, the script runs in the remote JVM or .Net Runtime.

Consumption of the ScriptOnConnector operation, at the API Level

```
@Test
public void runScriptOnConnectorTest() {
    logger.info("Running RunScriptOnConnector Test");
    final ConnectorFacade facade = createConnectorFacade(BasicConnector.class, null);
    final OperationOptionsBuilder builder = new OperationOptionsBuilder();
    builder.setRunAsUser("admin");
    builder.setRunWithPassword(new GuardedString("Passw0rd".toCharArray()));

    final ScriptContextBuilder scriptBuilder =
            new ScriptContextBuilder("Groovy", "return argument");
    scriptBuilder.addScriptArgument("argument", "value");

    Object result = facade.runScriptOnConnector(scriptBuilder.build(), builder.build());
    Assert.assertEquals(result, "value");
}
```

4.4.6.2. Implementing the Script on Connector Operation in Your Connector

The `scriptOnConnector` SPI operation takes the following parameters:

- `request` - the script and the arguments to be run

- `options` - additional options that control how the script is run

The operation returns the result of the script. The return type must be a type that the framework supports for serialization. See the ObjectSerializerFactory javadoc for a list of supported return types.

Implementation of the ScriptOnConnector operation, at the SPI Level

```java
public Object runScriptOnConnector(ScriptContext request, OperationOptions options) {
    final ScriptExecutorFactory factory =
            ScriptExecutorFactory.newInstance(request.getScriptLanguage());
    final ScriptExecutor executor =
            factory.newScriptExecutor(getClass().getClassLoader(), request.getScriptText(),
                    true);

    if (StringUtil.isNotBlank(options.getRunAsUser())) {
        String password = SecurityUtil.decrypt(options.getRunWithPassword());
        // Use these to execute the script with these credentials
    }
    try {
        return executor.execute(request.getScriptArguments());
    } catch (Throwable e) {
        logger.warn(e, "Failed to execute Script");
        throw ConnectorException.wrap(e);
    }
}
```

4.4.7. Script On Resource Operation

The script on resource operation runs a script directly on the target resource (unlike the "Script On Connector Operation", which runs a script in the context of a specific connector.)

Implement this interface if your connector intends to support the `ScriptOnResourceApiOp` API operation. If your connector implements this interface, you must document the script languages that the connector supports, as well as any supported `OperationOptions`.

4.4.7.1. Using the OpenICF Script on Resource Operation

The contract at the API level is intentionally very loose. Each connector decides what script languages it supports, what running a script on a target resource actually means, and what script options (if any) the connector supports.

Consumption of the ScriptOnResource operation, at the API Level

```
@Test
public void runScriptOnResourceTest() {
    logger.info("Running RunScriptOnResource Test");
    final ConnectorFacade facade = createConnectorFacade(BasicConnector.class, null);
    final OperationOptionsBuilder builder = new OperationOptionsBuilder();
    builder.setRunAsUser("admin");
    builder.setRunWithPassword(new GuardedString("Passw0rd".toCharArray()));

    final ScriptContextBuilder scriptBuilder = new ScriptContextBuilder("bash", "whoami");

    Object result = facade.runScriptOnResource(scriptBuilder.build(), builder.build());
    Assert.assertEquals(result, "admin");
}
```

4.4.7.2. Implementing the Script on Resource Operation in Your Connector

The `scriptOnResource` SPI operation takes the following parameters:

- `request` - the script and the arguments to be run

- `options` - additional options that control how the script is run

The operation returns the result of the script. The return type must be a type that the framework supports for serialization. See the ObjectSerializerFactory javadoc for a list of supported return types.

Implementation of the ScriptOnResource operation, at the SPI Level

```
public Object runScriptOnResource(ScriptContext request, OperationOptions options) {
    try {
        // Execute the script on remote resource
        if (StringUtil.isNotBlank(options.getRunAsUser())) {
            String password = SecurityUtil.decrypt(options.getRunWithPassword());
            // Use these to execute the script with these credentials
            return options.getRunAsUser();
        }
        throw new UnknownHostException("Failed to connect to remote SSH");
    } catch (Throwable e) {
        logger.warn(e, "Failed to execute Script");
        throw ConnectorException.wrap(e);
    }
}
```

4.4.8. Search Operation

The search operation enables the connector to search for objects on the target system.

The OpenICF framework handles searches as follows:

1. The application sends a query, with a search filter, to the OpenICF framework

2. The framework submits the query, with the filter, to the connector

3. The connector implements the `createFilterTranslator()` method to obtain a `FilterTranslator` object

4. The framework then uses this `FilterTranslator` object to transform the filter to a format that the `executeQuery()` method expects

You can implement the `FilterTranslator` object in two ways:

- The `FilterTranslator` translates the original filter into one or more native queries.

  The framework then calls the `executeQuery()` method for each native query.

- The `FilterTranslator` does not modify the original filter.

  The framework then calls the `executeQuery()` method with the original OpenICF filter.

  Using this second approach enables your connector to distinguish between a search and a get operation and to benefit from the visitor design pattern.

Based on the `resultsHandlerConfiguration`, the OpenICF framework can perform additional filtering on the returning results. For more information on the `resultsHandlerConfiguration`, see Results Handler Configuration.

The connector facade calls the `executeQuery` method once for each native query that the filter translator produces. If the filter translator produces more than one native query, the connector facade merges the results from each query and eliminates any duplicates.

Note that this implies an in-memory data structure that holds a set of UID values. Memory usage, in the event of multiple queries, will be O(N) where N is the number of results. It is therefore important that the filter translator for the connector implement `OR` operators, if possible.

Whether the application calls a `get` API operation, or a `search` API operation, the OpenICF framework translates that request to a `search` request on the connector.

4.4.8.1. Using the OpenICF Get Operation

The `GetApiOp` returns `null` when the UID does not exist on the resource.

Consumption of the Get operation, at the API Level

```java
@Test
public void getObjectTest() {
    logger.info("Running GetObject Test");
    final ConnectorFacade facade = createConnectorFacade(BasicConnector.class, null);
    final OperationOptionsBuilder builder = new OperationOptionsBuilder();
    builder.setAttributesToGet(Name.NAME);
    ConnectorObject co =
            facade.getObject(ObjectClass.ACCOUNT, new Uid(
                    "3f50eca0-f5e9-11e3-a3ac-0800200c9a66"), builder.build());
    Assert.assertEquals(co.getName().getNameValue(), "Foo");
}
```

4.4.8.2. Using the OpenICF Search Operation

Consumption of the Search operation, at the API Level

```java
@Test
public void searchTest() {
    logger.info("Running Search Test");
    final ConnectorFacade facade = createConnectorFacade(BasicConnector.class, null);
    final OperationOptionsBuilder builder = new OperationOptionsBuilder();
    builder.setPageSize(10);
    final ResultsHandler handler = new ToListResultsHandler();

    SearchResult result =
            facade.search(ObjectClass.ACCOUNT, FilterBuilder.equalTo(new Name("Foo")), handler,
                    builder.build());
    Assert.assertEquals(result.getPagedResultsCookie(), "0");
    Assert.assertEquals(((ToListResultsHandler) handler).getObjects().size(), 1);
}
```

4.4.8.3. Implementing the Search Operation in Your Connector

Implementation of the Search operation, at the SPI Level

```
public FilterTranslator<String> createFilterTranslator(ObjectClass objectClass,
        OperationOptions options) {
    return new BasicFilterTranslator();
}

public void executeQuery(ObjectClass objectClass, String query, ResultsHandler handler,
        OperationOptions options) {
    final ConnectorObjectBuilder builder = new ConnectorObjectBuilder();
    builder.setUid("3f50eca0-f5e9-11e3-a3ac-0800200c9a66");
    builder.setName("Foo");
    builder.addAttribute(AttributeBuilder.buildEnabled(true));

    for (ConnectorObject connectorObject : CollectionUtil.newSet(builder.build())) {
        if (!handler.handle(connectorObject)) {
            // Stop iterating because the handler stopped processing
            break;
        }
    }
    if (options.getPageSize() != null && 0 < options.getPageSize()) {
        logger.info("Paged Search was requested");
        ((SearchResultsHandler) handler).handleResult(new SearchResult("0", 0));
    }
}
```

4.4.9. Sync Operation

The sync operation polls the target system for synchronization events, that is, native changes to target objects.

The operation has two methods:

- `sync()` - request synchronization events from the target system

  This method calls the specified handler, once, to pass back each matching synchronization event. When the method returns, it will no longer invoke the specified handler.

- `getLatestSyncToken()` - returns the token corresponding to the most recent synchronization event

4.4.9.1. Using the OpenICF Sync Operation

Consumption of the Sync Operation (`getLatestSyncToken()` Method), at the API Level

```
@Test
public void getLatestSyncTokenTest() {
    logger.info("Running GetLatestSyncToken Test");
    final ConnectorFacade facade = createConnectorFacade(BasicConnector.class, null);
    SyncToken token = facade.getLatestSyncToken(ObjectClass.ACCOUNT);
    Assert.assertEquals(token.getValue(), 10);
}
```

The `getLatestSyncToken` method throws an `IllegalArgumentException` if the `objectClass` is null or invalid.

Consumption of the Sync Operation (`sync()` Method), at the API Level

```
@Test
public void syncTest() {
    logger.info("Running Sync Test");
    final ConnectorFacade facade = createConnectorFacade(BasicConnector.class, null);
    final OperationOptionsBuilder builder = new OperationOptionsBuilder();
    builder.setPageSize(10);
    final SyncResultsHandler handler = new SyncResultsHandler() {
        public boolean handle(SyncDelta delta) {
            return false;
        }
    };

    SyncToken token =
            facade.sync(ObjectClass.ACCOUNT, new SyncToken(10), handler, builder.build());
    Assert.assertEquals(token.getValue(), 10);
}
```

The `sync` method throws an `IllegalArgumentException` if the `objectClass` or `handler` is null, or if any argument is invalid.

4.4.9.2. Implementing the Sync Operation in Your Connector

Implementation of the Sync Operation at the SPI Level

```
public void sync(ObjectClass objectClass, SyncToken token, SyncResultsHandler handler,
        final OperationOptions options) {
    if (ObjectClass.ALL.equals(objectClass)) {
        //
    } else if (ObjectClass.ACCOUNT.equals(objectClass)) {
        final ConnectorObjectBuilder builder = new ConnectorObjectBuilder();
        builder.setUid("3f50eca0-f5e9-11e3-a3ac-0800200c9a66");
        builder.setName("Foo");
        builder.addAttribute(AttributeBuilder.buildEnabled(true));

        final SyncDeltaBuilder deltaBuilder = new SyncDeltaBuilder();
        deltaBuilder.setObject(builder.build());
```

```
        deltaBuilder.setDeltaType(SyncDeltaType.CREATE);
        deltaBuilder.setToken(new SyncToken(10));

        for (SyncDelta connectorObject : CollectionUtil.newSet(deltaBuilder.build())) {
            if (!handler.handle(connectorObject)) {
                // Stop iterating because the handler stopped processing
                break;
            }
        }
    } else {
        logger.warn("Sync of type {0} is not supported", configuration.getConnectorMessages()
                .format(objectClass.getDisplayNameKey(), objectClass.getObjectClassValue()));
        throw new UnsupportedOperationException("Sync of type"
                + objectClass.getObjectClassValue() + " is not supported");
    }
    ((SyncTokenResultsHandler) handler).handleResult(new SyncToken(10));
}

public SyncToken getLatestSyncToken(ObjectClass objectClass) {
    if (ObjectClass.ACCOUNT.equals(objectClass)) {
        return new SyncToken(10);
    } else {
        logger.warn("Sync of type {0} is not supported", configuration.getConnectorMessages()
                .format(objectClass.getDisplayNameKey(), objectClass.getObjectClassValue()));
        throw new UnsupportedOperationException("Sync of type"
                + objectClass.getObjectClassValue() + " is not supported");
    }
}
```

4.4.10. Test Operation

The test operation tests the connector configuration. Unlike validation, testing a configuration verifies that every part of the environment that is referred to by the configuration is available. The operation therefore validates that the connection details that are provided in the configuration are accurate, and that the backend is accessible when using them.

For example, the connector might make a physical connection to the host that is specified in the configuration, to check that it exists and that the credentials supplied in the configuration are valid.

The test operation can be invoked before the configuration has been validated, or can validate the configuration before testing it.

4.4.10.1. Using the OpenICF Test Operation

At the API level, the test operation throws a `RuntimeException` if the configuration is not valid, or if the test fails. Your connector implementation should throw the most specific exception available. When no specific exception is available, your connector implementation should throw a `ConnectorException`.

Consumption of the Test Operation at the API Level

```
@Test
public void testTest() {
    logger.info("Running Test Test");
    final ConnectorFacade facade = createConnectorFacade(BasicConnector.class, null);
    facade.test();
}
```

4.4.10.2. Implementing the Test Operation in Your Connector

Implementation of the Test Operation at the SPI Level

```
public void test() {
    logger.ok("Test works well");
}
```

4.4.11. Update Operation

If your connector will allow an authorized caller to update (modify or replace) objects on the target system, you must implement either the update operation, or the "Update Attribute Values Operation". At the API level update operation calls either the `UpdateOp` or the `UpdateAttributeValuesOp`, depending on what you have implemented.

The update operation is somewhat simpler to implement than the "Update Attribute Values Operation", because the update attribute values operation must handle any type of update that the caller might specify. However a true implementation of the update attribute values operation offers better performance and atomicity semantics.

4.4.11.1. Using the OpenICF Update Operation

At the API level, the update operation returns an `UnknownUidException` if the UID does not exist on the target system resource and if the connector does not implement the "Update Attribute Values Operation" interface.

Consumption of the Update Operation at the API Level

```
@Test
public void updateTest() {
    logger.info("Running Update Test");
    final ConnectorFacade facade = createConnectorFacade(BasicConnector.class, null);
    final OperationOptionsBuilder builder = new OperationOptionsBuilder();
    Set<Attribute> updateAttributes = new HashSet<Attribute>();
    updateAttributes.add(new Name("Foo"));

    Uid uid = facade.update(ObjectClass.ACCOUNT, new Uid("Foo"), updateAttributes, builder
                    .build());
    Assert.assertEquals(uid.getUidValue(), "foo");
}
```

4.4.11.2. Implementing the Update Operation in Your Connector

At the SPI level, the update operation returns an `UnknownUidException` if the UID does not exist on the target system.

Implementation of the Update Operation at the SPI Level

```
public Uid update(ObjectClass objectClass, Uid uid, Set<Attribute> replaceAttributes,
        OperationOptions options) {
    AttributesAccessor attributesAccessor = new AttributesAccessor(replaceAttributes);
    Name newName = attributesAccessor.getName();
    Uid uidAfterUpdate = uid;
    if (newName != null) {
        logger.info("Rename the object {0}:{1} to {2}", objectClass.getObjectClassValue(), uid
                .getUidValue(), newName.getNameValue());
        uidAfterUpdate = new Uid(newName.getNameValue().toLowerCase());
    }

    if (ObjectClass.ACCOUNT.equals(objectClass)) {

    } else if (ObjectClass.GROUP.is(objectClass.getObjectClassValue())) {
        if (attributesAccessor.hasAttribute("members")) {
            throw new InvalidAttributeValueException(
                    "Requested to update a read only attribute");
        }
    } else {
        logger.warn("Update of type {0} is not supported", configuration.getConnectorMessages()
                .format(objectClass.getDisplayNameKey(), objectClass.getObjectClassValue()));
        throw new UnsupportedOperationException("Update of type"
                + objectClass.getObjectClassValue() + " is not supported");
    }
    return uidAfterUpdate;
}
```

4.4.11.2.1. Suggested Approach for Deleting Attributes and Removing Attribute Values

If the target resource to which you are connecting supports the removal of attributes, you can implement the removal in several ways. All the samples in this document assume the following syntax rules for deleting attributes or removing their values.

Update	Syntax rule	Query filter
Set an empty attribute value	`[""]` (application sends an attribute value that is a list containing one empty string)	`equal=""`
Set an attribute value to null	`[]` (application sends an attribute value that is an empty list)	`ispresent` search returns 1
Removing an attribute	`null` (application sends an attribute value that is `null`	`ispresent` search returns 1

4.4.12. Update Attribute Values Operation

The update attribute values operation is an advanced implementation of the update operation. You should implement this operation if you want your connector to offer better performance and atomicity for the following methods:

- `UpdateApiOp.addAttributeValues(ObjectClass, Uid, Set, OperationOptions)`

- `UpdateApiOp.removeAttributeValues(ObjectClass, Uid, Set, OperationOptions)`

Consumption of the Add and Remove Attribute Values Methods at the API Level

```java
@Test
public void addAttributeValuesTest() {
    logger.info("Running AddAttributeValues Test");
    final ConnectorFacade facade = createConnectorFacade(BasicConnector.class, null);
    final OperationOptionsBuilder builder = new OperationOptionsBuilder();
    Set<Attribute> updateAttributes = new HashSet<Attribute>();
    // add 'group2' to existing groups
    updateAttributes.add(AttributeBuilder.build(PredefinedAttributes.GROUPS_NAME, "group2"));

    Uid uid =
            facade.addAttributeValues(ObjectClass.ACCOUNT, new Uid("Foo"), updateAttributes,
                    builder.build());
    Assert.assertEquals(uid.getUidValue(), "foo");
}

@Test
public void removeAttributeValuesTest() {
    logger.info("Running RemoveAttributeValues Test");
    final ConnectorFacade facade = createConnectorFacade(BasicConnector.class, null);
    final OperationOptionsBuilder builder = new OperationOptionsBuilder();
    Set<Attribute> updateAttributes = new HashSet<Attribute>();
    // remove 'group2' from existing groups
    updateAttributes.add(AttributeBuilder.build(PredefinedAttributes.GROUPS_NAME, "group2"));

    Uid uid =
            facade.removeAttributeValues(ObjectClass.ACCOUNT, new Uid("Foo"), updateAttributes,
                    builder.build());
    Assert.assertEquals(uid.getUidValue(), "foo");
}
```

4.4.12.1. Implementing the Update Attribute Values Operation in Your Connector

At the SPI level, the update attribute values operation returns an `UnknownUidException` when the UID does not exist on the resource.

Implementation of the update attribute values operation, at the SPI Level

```java
public Uid addAttributeValues(ObjectClass objectClass, Uid uid, Set<Attribute> valuesToAdd,
        OperationOptions options) {
    return uid;
}

public Uid removeAttributeValues(ObjectClass objectClass, Uid uid,
        Set<Attribute> valuesToRemove, OperationOptions options) {
    return uid;
}
```

4.5. Common Exceptions

The following sections describe the commonly used exceptions that can be thrown, depending on the operation.

4.5.1. AlreadyExistsException

The `AlreadyExistsException` is thrown if a create operation attempts to create an object that exists prior to the method execution, or if an update operation attempts to rename an object to that exists prior to the method execution.

4.5.2. ConfigurationException

A `ConfigurationException` is thrown if a configuration problem is encountered when the connector bundles are loaded. A `ConfigurationException` can also be thrown during validation operations in the SPI.

4.5.3. ConnectionBrokenException

A `ConnectionBrokenException` is thrown when a connection to a target resource instance fails during an operation. An instance of the `ConnectionBrokenException` generally wraps the native exception (or describes the native error) that is returned by the target resource.

4.5.4. ConnectionFailedException

A `ConnectionFailedException` is thrown when a connector cannot reach the target resource. An instance of the `ConnectionFailedException` generally wraps the native exception (or describes the native error) that is returned by the target resource.

4.5.5. ConnectorException

This is the base exception for the connector framework. The framework only throws exceptions that extend `ConnectorException`.

4.5.6. ConnectorIOException

This is the base exception for all Input-Output (I/O-related) exceptions, including instance connection failure, socket error and so forth.

4.5.7. ConnectorSecurityException

This is the base exception for all security-related exceptions.

4.5.8. InvalidAttributeValueException

An `InvalidAttributeValueException` is thrown when an attempt is made to add to an attribute a value that conflicts with the attribute's schema definition. This might happen, for example, in the following situations:

- The connector attempts to add an attribute with no value when the attribute is required to have at least one value

- The connector attempts to add more than one value to a single valued-attribute

- The connector attempts to add a value that conflicts with the attribute type

- The connector attempts to add a value that conflicts with the attribute syntax

4.5.9. InvalidCredentialException

An `InvalidCredentialException` indicates that user authentication has failed. This exception is thrown by the connector when authentication fails, and when the specific reason for the failure is not known. For example, the connector might throw this exception if a user has entered an incorrect password, or username.

4.5.10. InvalidPasswordException

An `InvalidPasswordException` is thrown when a password credential is invalid.

4.5.11. OperationTimeoutException

An `OperationTimeoutException` is thrown when an operation times out. The framework cancels an operation when the corresponding method has been executing for longer than the limit specified in `APIConfiguration`.

4.5.12. PasswordExpiredException

A `PasswordExpiredException` indicates that a user password has expired. This exception is thrown by the connector when it can determine that a password has expired. For example, after successfully authenticating a user, the connector might determine that the user's password has expired. The connector throws this exception to notify the application, which can then take the appropriate steps to notify the user.

4.5.13. PermissionDeniedException

A `PermissionDeniedException` is thrown when the target resource will not allow a connector to perform a particular operation. An instance of the `PermissionDeniedException` generally describes a native error (or wraps a native exception) that is returned by the target resource.

4.5.14. PreconditionFailedException

A `PreconditionFailedException` is thrown to indicate that a resource's current version does not match the version provided. This exception is equivalent to the HTTP status: `412 Precondition Failed`.

4.5.15. PreconditionRequiredException

A `PreconditionRequiredException` is thrown to indicate that a resource requires a version, but that no version was supplied in the request. This exception is equivalent to the HTTP status: `428 Precondition Required`.

4.5.16. RetryableException

A `RetryableException` indicates that the failure might be temporary, and that retrying the same request might succeed in the future.

4.5.17. UnknownUidException

An `UnknownUidException` is thrown when a UID that is specified as input to a connector operation identifies no object on the target resource. When you implement the `AuthenticateOp`, your connector can throw this exception if it is unable to locate the account necessary to perform authentication.

4.5.18. NullPointerException (c# NullReferenceException)

Generic native exception

4.5.19. UnsupportedOperationException (c# NotSupportedException)

Generic native exception

4.5.20. IllegalStateException (c# InvalidOperationException)

Generic native exception

4.5.21. IllegalArgumentException (c# ArgumentException)

Generic native exception

4.5.22. Mapping OpenICF Exceptions to ForgeRock CREST Exceptions

The following table maps the errors that are thrown by the OpenICF framework to the errors that are returned by the ForgeRock Commons REST (CREST) implementation.

OpenICF Exception	CREST Exception	HTTP Error Code
AlreadyExistsException	ConflictException	
ConfigurationException	InternalServerErrorException	
ConnectionBrokenException	InternalServerErrorException	
ConnectionFailedException	ConnectionFailedException	
ConnectorException	InternalServerErrorException	
ConnectorIOException	InternalServerErrorException	
ConnectorSecurityException	ForbiddenException	
InvalidAttributeValueException	BadRequestException	
InvalidCredentialException	ForbiddenException	
InvalidPasswordException	ForbiddenException	
OperationTimeoutException		
PasswordExpiredException	ForbiddenException	
PermissionDeniedException	ForbiddenException	
PreconditionFailedException	PreconditionFailedException	
PreconditionRequiredException	PreconditionRequiredException	
RetryableException	RetryableException (ServiceUnavailableException)	
UnknownUidException	NotFoundException	
UnsupportedOperationException	NotSupportedException	
IllegalArgumentException	InternalServerErrorException	
NullPointerException	InternalServerErrorException	

4.6. Generic Exception Rules

The generic exception rules are common to all API or SPI level operations and are described in the following sections.

4.6.1. Framework (API Level) Exception Rules

IllegalArgumentException or NullPointerException

> Thrown when the ObjectClass is null or the name is blank.

OperationTimeoutException

> Thrown when the operation timed out.

ConnectionFailedException

Thrown if any problem occurs with the connector server connection.

UnsupportedOperationException

Thrown if the connector does not implement the required interface.

ConnectorIOException

Thrown if the connector failed to initialize a remote connection due to a SocketException.

ConnectorException

Thrown in the following situations:

- The connector failed to initiate the remote connection due to a `SocketException`

- An unexpected request was sent to the remote connector server

- An unexpected response was received from the remote connector server

InvalidCredentialException

Thrown if the remote framework key is invalid

The following exceptions are thrown specifically in the context of a poolable connector.

ConnectorException

Thrown if the pool has no available connectors after the `maxWait` time has elapsed.

IllegalStateException

Thrown if the object pool has already shut down.

4.6.2. Connector (SPI Level) Exception Rules

InvalidAttributeValueException

Thrown when single-valued attribute has multiple values.

IllegalArgumentException

Thrown when the value of the `__PASSWORD__` or the `__CURRENT_PASSWORD__` attribute is not a `GuardedString`.

IllegalStateException

Thrown when the `Attribute` name is blank.

PermissionDeniedException

Thrown when the target resource will not allow a specific operation to be performed. An instance of the `PermissionDeniedException` generally describes a native error that is returned by (or wraps a native exception that is thrown by) the target resource.

ConnectorIOException, ConnectionBrokenException, ConnectionFailedException

Thrown when any problem occurs with the connection to the target resource.

PreconditionFailedException

Thrown when the current version of the resource object does not match the version provided by the connector.

PreconditionRequiredException

Thrown when a resource object requires a version, but no version was supplied in the `getRevision` operation.

Chapter 5
Writing Java Connectors

If none of the existing OpenICF connectors are suitable for your deployment, you can write your own connector. This chapter describes the steps to develop an OpenICF-compatible Java connector. Similar chapters exist to help you with writing scripted Groovy, and PowerShell connectors.

5.1. Deciding What Kind of Connector to Write

In general, it takes longer to write a new Java connector than it does to use one of the scripted connector toolkits to write a scripted connector. Before you can write a Java connector, you must have a good understanding of the OpenICF SPI (see "*Implementing the OpenICF SPI*").

Scripted connectors do not require a complete understanding of the SPI, so connector development should be faster. The scripted connector implementations provided with IDM follow a general *pattern* and you can assess which implementation to use based on what the connector must be able to do.

For example, if you need to connect to a database, use the scriptedSQL implementation. To execute a remote command over SSH, use the scriptedSSH implementation. The details of these different scripted connector types are described in "Selecting a Scripted Connector Implementation".

If the main purpose of your connector is to call a number of stored procedures or perform some SQL inserts, you can avoid learning the OpenICF SPI and focus on the required "actions" (create, delete, update, and so on). You can then implement these actions in a scripted connector. When you have stable scripts that do what they need to do, package them in a .jar, version them and your connector development is complete.

If you need to connect to *new* system with a client/server API in written in Java, you must write a new Java connector. This chapter helps you get started with that process.

5.2. Before You Begin

Before you start developing your own connector, familiarize yourself with the structure of the SPI, by reading "*Implementing the OpenICF SPI*" and the corresponding Javadoc for the OpenICF framework and its supported operations.

5.3. Using the Connector Archetype

OpenICF provides a Maven connector archetype that enables you to get started with connector development.

The connector archetype assumes that you have Apache Maven installed on your system. Before you use the connector archetype, add the following to your Maven `settings.xml` file, replacing `backstage-username` and `backstage-password` with your ForgeRock Backstage credentials:

```xml
<servers>
 ...
  <server>
    <username>backstage-username</username>
    <password>backstage-password</password>
    <id>archetype</id>
  </server>
</servers>
...
<profiles>
  <profile>
    <id>test</id>
    <activation>
      <activeByDefault>true</activeByDefault>
    </activation>
    <repositories>
      <repository>
          <id>archetype</id>
          <url>http://maven.forgerock.org/repo/private-releases</url>
        </repository>
    </repositories>
  </profile>
</profiles>
```

To start building a connector by using the connector archetype, execute the following command, customizing these options to describe your new connector:

- `-DartifactId=sample-connector`

- `-Dversion=0.0-SNAPSHOT`

- `-Dpackage=org.forgerock.openicf.connectors.sample`

- `-DconnectorName=Sample`

This command imports the connector archetype and generates a new connector project:

```
$ mvn archetype:generate \
    -DarchetypeGroupId=org.forgerock.openicf \
    -DarchetypeArtifactId=connector-archetype \
    -DarchetypeVersion=1.4.0 \
    -DremoteRepositories=http://maven.forgerock.org/repo/private-releases \
    -DarchetypeRepository=http://maven.forgerock.org/repo/private-releases \
    -DgroupId=org.forgerock.openicf.connectors \
    -DartifactId=sample-connector \
    -Dversion=0.0-SNAPSHOT \
```

```
      -Dpackage=org.forgerock.openicf.connectors.sample \
      -DconnectorName=Sample
[INFO] Scanning for projects...
[INFO]
[INFO] ------------------------------------------------------------------------
[INFO] Building Maven Stub Project (No POM) 1
[INFO] ------------------------------------------------------------------------
[INFO]
[INFO] >>> maven-archetype-plugin:3.0.1:generate (default-cli) > generate-sources @ standalone-pom >>>
[INFO]
[INFO] <<< maven-archetype-plugin:3.0.1:generate (default-cli) < generate-sources @ standalone-pom <<<
[INFO]
[INFO] --- maven-archetype-plugin:3.0.1:generate (default-cli) @ standalone-pom ---
[INFO] Generating project in Interactive mode
...
ALL_OPERATIONS: n
OP_AUTHENTICATE: n
OP_CREATE: y
OP_DELETE: y
OP_RESOLVEUSERNAME: n
OP_SCHEMA: n
OP_SCRIPTONCONNECTOR: n
OP_SCRIPTONRESOURCE: n
OP_SEARCH: y
OP_SYNC: n
OP_TEST: y
OP_UPDATE: y
OP_UPDATEATTRIBUTEVALUES: n
attributeNormalizer: n
compatibility_version: 1.1
connectorName: Sample
framework_version: 1.0
jira_componentId: 10191
jira_fixVersionIds: 0
poolableConnector: n
 Y: :
```

At this point, you can enter Y (YES) to accept the default project, or N (NO) to customize the project for your connector.

You will notice in the preceding output that the default connector supports only the `create`, `delete`, `search`, `test`, and `update` operations, and is not a poolable connector. To add support for additional operations, or to change any of the connector parameters, enter N (NO). The archetype then prompts you to set values for each additional parameter.

After you have imported the archetype once, you can use the local version of the archetype, as follows:

```
$ mvn archetype:generate -DarchetypeCatalog=local
```

5.4. Implementing OpenICF Operations

When you have generated the archetype, implement the OpenICF operations that your connector will support.

For information about implementing operations, and examples for a Java connector, see "*Implementing the OpenICF SPI*".

Then build the connector, as follows:

5.5. Building the Java Connector

```
$ cd sample-connector
$ mvn install
```

Chapter 6

Writing Scripted Connectors With the Groovy Connector Toolkit

The Groovy Connector Toolkit enables you to run Groovy scripts to interact with any external resource.

The Groovy Connector Toolkit is not a complete connector, in the traditional sense. Rather, it is a framework within which you must write your own Groovy scripts to address the requirements of your deployment. The toolkit is bundled with IDM in the JAR `openidm/connectors/groovy-connector-1.5.20.0.jar`.

IDM provides a number of deployment-specific scripts to help you get started with the Groovy Connector Toolkit. These scripts demonstrate how the toolkit can be used. The scripts cannot be used "as is" in your deployment, but can be used as a starting point on which to base your customization.

The Groovy Connector Toolkit can be used with any OpenICF-enabled project (that is, any project in which the OpenICF Connector Framework is installed).

6.1. About the Groovy Scripting Language

Groovy is a powerful, convenient scripting language for the Java platform. Groovy enables you to take advantage of existing Java resources, and generally makes development quicker. Syntactically, Groovy is similar to JavaScript. Extensive information about Groovy is available on the Groovy documentation site.

6.2. Selecting a Scripted Connector Implementation

The Groovy Connector Toolkit provides five default connector implementations. The default implementations should address the requirements of most target resources. If you use one of the default implementations, you need only write the accompanying scripts and point your connector to their location. If your target resource is not covered by the default implementations, you can use the Maven archetype to create a new connector project, and write a custom Groovy-based connector from scratch.

The following list describes the default scripted connector implementations provided with the Groovy Connector Toolkit:

• `GROOVY` - a basic non-pooled Groovy connector, provided in the `org.forgerock.openicf.connectors.groovy.ScriptedConnector` class.

`POOLABLEGROOVY` - a poolable Groovy connector, provided in the `org.forgerock.openicf.connectors.groovy` `.ScriptedPoolableConnector` class.

`REST` - a scripted REST connector, provided in the `org.forgerock.openicf.connectors.groovy` `.ScriptedRESTConnector` class. The scripted REST connector enables you to connect to any resource, over HTTP/REST. The connector creates the HTTP/REST context (specifying the content type, authentication mode, encoding, and so on), and manages the connection. The connector relies on the Groovy scripting language and its RESTClient package.

`SQL` - a scripted SQL connector, provided in the `org.forgerock.openicf.connectors.groovy` `.ScriptedSQLConnector` class. The scripted SQL connector uses Groovy scripts to interact with a JDBC database.

When you have selected a scripted connector implementation, write the required scripts that correspond to that connector type. "Implementing OpenICF Operations With Groovy Scripts" provides information and examples on how to write scripts for the basic scripted connector implementation, and information on the extensions available for the other implementations.

6.3. Starting With the Groovy Sample Project

ForgeRock provides a sample project that you can use to start writing a custom scripted Groovy connector. The sample project includes the base classes and libraries required for a simple connector, but the built connector does nothing until you add scripts to perform the operations on your remote resource.

The main use of the sample project is to make it easier to package your custom Groovy scripts into a single custom scripted connector file.

To get started with the sample project:

1. Clone the project from the ForgeRock git repository.

    If you do not have access to this repository, contact info@forgerock.com.

2. Open the project in your IDE.

3. Rename the connector.

    In general, you can find and replace all instances of `customscripted` in your project's `pom.xml` file with the name of your connector. The following examples assume a new connector named `myexample`:

    a. Rename `org.forgerock.openicf.connectors.customscripted` to `org.forgerock.openicf.connectors` `.myexample`. Specify your own directory structure here.

    b. In the `pom.xml` file, make the following changes:

        • Change the `artifactId`:

```
<artifactId>myexample-connector</artifactId>
```

- Change the description as required. For example:

```
<description>Connector that manages users in the myexample resource</description>
```

- Change the `connectorPackage` and `connectorClass`:

```
<connectorPackage>org.forgerock.openicf.connectors.myexample</connectorPackage>
<connectorClass>MyexampleConnector</connectorClass>
```

- Edit the `repositories` block to include additional repositories if necessary.

  At a minimum, the connector pulls the `connector-framework` and `groovy-connector` artifacts from the `forgerock-private-releases` repository.

- In the `maven-bundle-plugin` configuration, change the export package name:

```
<Export-Package>org.forgerock.openicf.connectors.myexample.*</Export-Package>
```

4. Edit the display messages in the `resources/../Messages.properties` file.

   The `connector.display` property specifies the name of this connector as it will appear in the UI.

   Add a `.display` and `.help` property for any configuration properties for the connector. These properties are specific to each connector. For example, if your connector connects to a database, you might need to specify the properties similar to the following here:

```
database.display=Database Path
database.help=If specified, the connector attempts to connect using only this database
```

   The `Messages.properties` file is not mandatory, and is used only to manage the connector configuration in the UI.

5. Rename the `CustomscriptedConnector` class. For example:

```
@ConnectorClass(displayNameKey = "connector.display",
        configurationClass = MyexampleConfiguration.class,
        messageCatalogPaths = {
                "org/forgerock/openicf/connectors/groovy/Messages",
                "org/forgerock/openicf/connectors/customscripted/Messages"
        })
public class Myexample extends ScriptedConnectorBase<MyexampleConfiguration>
        implements PoolableConnector {

    public Myexample() {
        // Empty
    }
```

   In this example, your connector class extends the `ScriptedConnectorBase` class. You can also extend the, `ScriptedREST`, `ScriptedSQL`, or `ScriptedSSH` class, with additional imports, if your connector should

inherit those capabilities. If you extend a different connector class here, you must also extend the corresponding `Configuration` class in the `CustomscriptedConfiguration` shown in the next step.

6. Add any new methods or override the default behavior in this `Connector` class, as required.

7. Rename the `CustomscriptedConfiguration` class. For example:

```
@ConfigurationClass(skipUnsupported = true)
public class MyexampleConfiguration extends ScriptedConfiguration {

    /**
     * Setup logging for the {@link MyexampleConfiguration}.
     */
    private static final Log logger = Log.getLog(MyexampleConfiguration.class);

    private String example = null;
    private GuardedString guardedExample = null;

    /**
     * Constructor.
     */
    public MyexampleConfiguration() {
    }
```

8. Add any new configuration properties that your connector requires (with getters, setters, and annotations).

9. Update your project's `pom.xml` file to include any new dependencies.

   If you do not want to bundle the dependencies with the connector, set `<scope>provided</scope>` for those dependencies. For example:

```
<dependency>
    <groupId>org.codehaus.groovy</groupId>
    <artifactId>groovy-all</artifactId>
    <version>2.2.2</version>
    <scope>provided</scope>
</dependency>
```

10. Write the scripts that will perform the OpenICF operations on the resource, and place these scripts in the `src/main/resources/script/` directory of the project.

    IDM provides several sample scripts in its scripted connector samples and you can use these as a base from which to work. See the scripts in the following sample directories to get started:

    * `/path/to/openidm/samples/scripted-rest-with-dj/tools`

    * `/path/to/openidm/samples/scripted-sql-with-mysql/tools`

    To use the sample scripts, copy them from the sample directory and place them your project's `script/` directory, then edit them as required. For example:

```
$ cd /path/to/myexample-connector/
$ cp /path/to/openidm/samples/scripted-sql-with-mysql/tools/TestScript.groovy src/main/resources/
script/
```

11. Build your connector with the following command:

```
$ mvn clean install
```

When your connector has built successfully, the `target/` directory contains two JAR files:

- `myexample-connector-1.5.0.0-SNAPSHOT.jar` bundles only the connector and its scripts.

- `myexample-connector-1.5.0.0-SNAPSHOT-jar-with-dependencies.jar` bundles the connector with its scripts and dependencies.

> **Note**
>
> If you want to build a single JAR file, including the dependencies, named `myexample-connector-1.5.0.0-SNAPSHOT.jar`, edit your `pom.xml` accordingly.

12. To use your connector, create a connector configuration (provisioner) file using one of the methods described in "Configuring Connectors" in the *Integrator's Guide*.

13. Edit the provisioner file as required. At a minimum, set the `scriptRoots` property to point to the location of your scripts within the bundled connector JAR. For example:

```
"scriptRoots" : [ "jar:file:connectors/myexample-connector-1.5.0.0-SNAPSHOT.jar!/script/" ]
```

6.4. Implementing OpenICF Operations With Groovy Scripts

The Groovy Connector Toolkit enables you to run a Groovy script for any OpenICF operation, such as search, update, create, and so forth, on any external resource.

You must write a Groovy script that corresponds to each operation that your connector will support. For information about all the operations that are supported by the OpenICF framework, see "*Implementing the OpenICF SPI*".

Your scripted connector can implement the following OpenICF interfaces:

"Authenticate Operation"

Provides simple authentication with two parameters, presumed to be a user name and password.

"Create Operation"

Creates an object and its `uid`.

"Delete Operation"

Deletes an object, referenced by its `uid`.

"Resolve Username Operation"

Resolves an object to its `uid` based on its `username`.

"Schema Operation"

Describes the object types, operations, and options that the connector supports.

"Script On Connector Operation"

Enables IDM to run a script in the context of the connector. Any script that runs on the connector has the following characteristics:

- The script runs in the same execution environment as the connector and has access to all the classes to which the connector has access.

- The script has access to a connector variable that is equivalent to an initialized instance of the connector. At a minimum, the script can access the connector configuration.

- The script has access to any script-arguments passed in by IDM.

"Script On Resource Operation"

Runs a script directly on the target resource that is managed by the connector.

"Search Operation"

Searches the target resource for all objects that match the specified object class and filter.

"Sync Operation"

Polls the target resource for synchronization events, that is, native changes to objects on the target resource.

"Test Operation"

Tests the connector configuration. Testing a configuration checks that all elements of the environment that are referred to by the configuration are available. For example, the connector might make a physical connection to a host that is specified in the configuration to verify that it exists and that the credentials that are specified in the configuration are valid.

This operation might need to connect to the resource, and, as such, might take some time. Do not invoke this operation too often, such as before every provisioning operation. The test operation is not intended to check that the connector is alive (that is, that its physical connection to the resource has not timed out).

You can invoke the test operation before a connector configuration has been validated.

"Update Operation"

Updates (modifies or replaces) objects on a target resource.

The following sections provide more information and pointers to sample scripts for all the operations that are implemented in the Groovy Connector Toolkit.

6.4.1. Variables Available to All Groovy Scripts

The following variables are available to all scripts used by the Groovy Connector. Additional variables are available to specific scripts, as described in the sections that follow:

`configuration`

A handle to the connector's configuration object is injected into all scripts.

`operation`

The connector injects the name of the action or operation into the script, to indicate which action is being called.

The sample scripts for the Groovy connector define one script file per action. You can use a single file, or amalgamate multiple actions into one file. For example, the CREATE and UPDATE operations often share the same code.

The operation type can be one of the following:

- `ADD_ATTRIBUTE_VALUES`

- `AUTHENTICATE`

- `CREATE`

- `DELETE`

- `GET_LATEST_SYNC_TOKEN`

- `REMOVE_ATTRIBUTE_VALUES`

- `RESOLVE_USERNAME`

- `RUNSCRIPTONCONNECTOR`

- `RUNSCRIPTONRESOURCE`

- `SCHEMA`

- `SEARCH`

- `SYNC`

- `TEST`

- UPDATE

options

The OpenICF framework passes an `OperationOptions` object to most of the operations. The Groovy connector injects this object, as is, into the scripts. For example, the search, query, and sync operations pass the attributes to get as an operation option.

The most common options are as follows:

- `AttributesToGet` (String[]) for search and sync operations

- `RunAsUser` (String) for any operation

- `RunWithPassword` (GuardedString) for any operation

- `PagedResultsCookie` (String) for search operations

- `PagedResultsOffset` (Int) for search operations

- `PageSize` (Int) for search operations

- `SortKeys` (Sortkey[]) for search operations

objectClass

The category or type of object that is managed by the connector, such as ACCOUNT and GROUP.

log

A handle to the default OpenICF logging facility.

connection

Available to the ScriptedREST, and ScriptedSQL implementations, this variable initiates the HTTP or SQL connection to the resource.

6.4.2. Writing an Authenticate Script

An authenticate script is *required* if you want to use pass-through authentication to the backend resource. If your connector does not need to authenticate to the resource, the authenticate script should allow the `authId` to pass through by default.

A sample authenticate script for an SQL database is provided in `openidm/samples/scripted-sql-with-mysql/tools/AuthenticateScript.groovy`

Input variables:

The following variables are available to the authenticate script:

configuration

> A handler to the connector's configuration object.

options

> A handler to the Operation Options.

operation

> An OperationType that corresponds to the action (`AUTHENTICATE`).

objectClass

> The object class being used to authenticate, such as `__ACCOUNT__` or `__GROUP__`.

username

> A string that provides the username to authenticate.

password

> A guarded string that provides the password with which to authenticate.

log

> A logger instance for the connector.

Returns: The user unique ID (OpenICF `__UID__`). The `type` of the returned UID must be a `string` or a `Uid`. The script must throw an exception in the case of failure.

Authenticate Script

```
def operation = operation as OperationType
def configuration = configuration as ScriptedConfiguration
def username = username as String
def log = log as Log
def objectClass = objectClass as ObjectClass
def options = options as OperationOptions
def password = password as GuardedString;

if (username.equals("TEST")) {
    def clearPassword = SecurityUtil.decrypt(password)
    if ("Passw0rd".equals(clearPassword)) {
        return new Uid(username);
    }
}
```

6.4.3. Writing a Test Script

A test script tests the connection to the external resource to ensure that the other operations that are provided by the connector can succeed.

A sample test script for an SQL database is provided in `openidm/samples/scripted-sql-with-mysql/tools/TestScript.groovy`

Input variables:

The following variables are available to the test script:

configuration

A handler to the connector's configuration object.

operation

An OperationType that corresponds to the action (`TEST`).

log

A logger instance for the connector.

Returns: Nothing, if the test is successful. The script can throw any exception if it fails.

Test Script

```
import org.identityconnectors.common.logging.Log
 import org.forgerock.openicf.connectors.groovy.OperationType
 import org.forgerock.openicf.misc.scriptedcommon.ScriptedConfiguration

 def operation = operation as OperationType
 def configuration = configuration as ScriptedConfiguration
 def log = log as Log

 log.info("This is a TestScript")
 throw new MissingResourceException("Test Failed", operation.name(), "")
```

6.4.4. Writing a Create Script

A create script creates a new object on the external resource. If your connector does not support creating an object, this script should throw an `UnsupportedOperationException`.

A sample create script for an SQL database is provided in `openidm/samples/scripted-sql-with-mysql/tools/CreateScript.groovy`

Input variables:

The following variables are available to a create script:

configuration

A handler to the connector's configuration object.

options

> A handler to the Operation Options.

operation

> An OperationType that corresponds to the action (`CREATE`).

objectClass

> The object class that is created, such as `__ACCOUNT__` or `__GROUP__`.

attributes

> The set of attributes that describe the object to be created.

id

> The UID of the object to be created, if specified. If the UID is `null`, the UID should be generated by the server. The UID corresponds to the OpenICF `__NAME__` attribute if it is provided as part of the attribute set.

log

> A logger instance for the connector.

Returns: The user unique ID (OpenICF `__UID__`) of the newly created object. The `type` of the returned UID must be a `string` or a `Uid`. If a null value or an object type other than `string` or `Uid` is returned, the script must throw an exception.

Create Script

```
def operation = operation as OperationType
def configuration = configuration as SapConfiguration
def log = log as Log
def objectClass = objectClass as ObjectClass
def createAttributes = new AttributesAccessor(attributes as Set<Attribute>)
def name = id as String
def options = options as OperationOptions

log.info("Entering {0} script",operation);


assert operation == OperationType.CREATE, 'Operation must be a CREATE'
// We only deal with users
assert objectClass.getObjectClassValue() == ObjectClass.ACCOUNT_NAME


def password = createAttributes.getPassword() as GuardedString;
assert password != null, 'Password must be provided on create'


//...
def uid = createTheUser(createAttributes);
return uid
```

6.4.5. Writing a Search or Query Script

A search script searches for one or more objects on the external resource. Connectors that do not support searches should throw an `UnsupportedOperationException`.

A sample search script for an SQL database is provided in `openidm/samples/scripted-sql-with-mysql/tools/SearchScript.groovy`

Input variables:

The following variables are available to the search script:

configuration

> A handler to the connector's configuration object.

options

> A handler to the Operation Options.

operation

> An OperationType that corresponds to the action (`SEARCH`).

objectClass

> The object class to search, such as `__ACCOUNT__` or `__GROUP__`.

filter

> The OpenICF native Query filter for this operation.

query

> A Map representation of the native Query filter that is easy to process.

> Provides a convenient way to access the query filter parameter. For example:

```
query = [ operation: "CONTAINS", left: attribute, right: "value", not: true/false ]
query = [ operation: "ENDSWITH", left: attribute, right: "value", not: true/false ]
query = [ operation: "STARTSWITH", left: attribute, right: "value", not: true/false ]
query = [ operation: "EQUALS", left: attribute, right: "value", not: true/false ]
query = [ operation: "GREATERTHAN", left: attribute, right: "value", not: true/false ]
query = [ operation: "GREATERTHANOREQUAL", left: attribute, right: "value", not: true/false ]
query = [ operation: "LESSTHAN", left: attribute, right: "value", not: true/false ]
query = [ operation: "LESSTHANOREQUAL", left: attribute, right: "value", not: true/false ]
query = null : then we assume we fetch everything

// AND and OR filter - embed these left/right queries:
query = [ operation: "AND", left: query1, right: query2 ]
query = [ operation: "OR", left: query1, right: query2 ]
```

For example, the equality query filter `"sn == Smith"` would be represented by the following query Map:

```
[ operation: "EQUALS", left: "sn", right: "Smith", not: false ]
```

handler

A Closure handler for processing the search results.

log

A logger instance for the connector.

Returns: Optionally, the script can return a search result. The result can be be returned as a `SearchResult` object or as a `String` that represents the `pagedResultsCookie` to be used for the next paged results.

6.4.5.1. Returning Search Results

In a search operation, a result handler (callback) is passed to the script to return the results one by one. The handler must be called for every query result. The handler variable that is passed to the script is a `Groovy Closure`. You can call the handler in the following ways:

- Using an OpenICF `ConnectorObject` object.

  You can use the `ConnectorObjectBuilder` to build this object. For example:

```
def builder = new ConnectorObjectBuilder()
  builder.setUid(uidValue)
  builder.setName(nameValue)
  builder.setObjectClass(ObjectClass.ACCOUNT)
  builder.addAttribute("sn", snValue)

  // Call the handler with the ConnectorObject object
  handler builder.build()
```

- Using a Groovy Closure.

  In this case the Closure delegates calls to a specific Object that implements these calls. For example:

```
handler {            // The handler parameter here is a Closure
    uid uidValue     // (mandatory), the method resolution for 'uid' is delegated to the Object
                        handling the Closure. This is the OpenICF __UID__
    id nameValue     // (mandatory), the method resolution for 'id' is delegated to the Object
                        handling the Closure. This is the OpenICF __NAME__
    attribute "sn", snValue // (optional), the method resolution for 'id' is delegated to the
                              Object handling the Closure
    // attribute <attribute2Name>, <attribute2Value>
    // etc...
}
```

In the following example, the handler is called within a loop to return all the results of a query:

```
for (user in userList) {
    handler {
                uid user.userName
                id user.userName
                user.attributes.each(){ key,value -> attribute key, value }
            }
}
```

6.4.6. Writing an Update Script

An update script updates an object in the external resource. Connectors that do not support update operations should throw an `UnsupportedOperationException`.

A sample update script for an SQL database is provided in `openidm/samples/scripted-sql-with-mysql/tools/UpdateScript.groovy`

Input variables:

The following variables are available to an update script:

configuration

A handler to the connector's configuration object.

options

A handler to the Operation Options.

operation

An OperationType that corresponds to the action (`UPDATE`).

objectClass

The object class that is updated, such as `__ACCOUNT__` or `__GROUP__`.

attributes

A collection of `ConnectorAttributes` that represent the entry attributes to update.

uid

The UID of the object to be updated. The UID corresponds to the OpenICF `UID` attribute.

id

The name of the object to be updated (optional). The id corresponds to the OpenICF `__NAME__` attribute. It will not be injected and set unless the update is a rename.

log

A logger instance for the connector.

Returns: The user unique ID (OpenICF `__UID__`) of the updated object. The `type` of the returned UID must be a `string` or a `Uid`. If the UID is not modified by the update operation, return the value of the uid injected into the script.

Update Script

```
def operation = operation as OperationType
def updateAttributes = attributes as Set<Attribute>
def configuration = configuration as ScriptedConfiguration
def id = id as String
def log = log as Log
def objectClass = objectClass as ObjectClass
def options = options as OperationOptions
def uid = uid as Uid

log.ok("Update...")
switch (operation) {
    case OperationType.UPDATE:
        switch (objectClass) {
            case ObjectClass.ACCOUNT:
// ...
                for (Attribute a : updateAttributes) {
                    if (a.is(Name.NAME)) {
// ...
return uid
```

6.4.7. Writing a Delete Script

A delete script deletes an object in the external resource. Connectors that do not support delete operations should throw an `UnsupportedOperationException`.

A sample delete script for an SQL database is provided in `openidm/samples/scripted-sql-with-mysql/tools/DeleteScript.groovy`

Input variables:

The following variables are available to an update script:

configuration

A handler to the connector's configuration object.

options

A handler to the Operation Options.

operation

An OperationType that corresponds to the action (`DELETE`).

objectClass

The object class that is deleted, such as `__ACCOUNT__` or `__GROUP__`.

uid

> The UID of the object to be deleted. The UID corresponds to the OpenICF `__UID__` attribute.

log

> A logger instance for the connector.

Returns: This script has no return value but should throw an exception if the delete is unsuccessful.

6.4.8. Writing a Synchronization Script

A synchronization script synchronizes objects between two resources. The script should retrieve all objects in the external resource that have been updated since some defined token.

A sample synchronization script for an SQL database is provided in `openidm/samples/scripted-sql-with-mysql/tools/SyncScript.groovy`

Input variables:

The following variables are available to a sync script:

configuration

> A handler to the connector's configuration object.

options

> A handler to the Operation Options.

operation

> An OperationType that corresponds to the action (`GET_LATEST_SYNC_TOKEN` or `SYNC`).

objectClass

> The object class that is synchronized, such as `__ACCOUNT__` or `__GROUP__`.

token

> The value of the sync token.

handler

> A `Closure` handler for processing the sync results.

log

> A logger instance for the connector.

Returns:

If the operation type is `GET_LATEST_SYNC_TOKEN`, the script must return an object that represents the last known `SyncToken` for the corresponding `ObjectClass`. For example:

```
def operation = operation as OperationType
def configuration = configuration as ScriptedConfiguration
def log = log as Log
def objectClass = objectClass as ObjectClass
def options = options as OperationOptions
def token = token as Object

case OperationType.GET_LATEST_SYNC_TOKEN:
        switch (objectClass) {
            case ObjectClass.ACCOUNT:
                return new SyncToken(17);
            case ObjectClass.GROUP:
                return new SyncToken(16);
            case ObjectClass.ALL:
                return new SyncToken(17);
// ....
```

If the operation type is `SYNC`, the script must return a new `SyncToken` for the corresponding `ObjectClass`. A Sync result handler (callback) is passed to the script to return the Sync results one by one. The handler must be called for each result.

The handler variable that is passed to the script is a Groovy Closure. It can be called in the following ways:

- With an OpenICF `SyncDelta` object.

  You can use a `SyncDeltaBuilder` to build this object. For example:

```
def builder = new SyncDeltaBuilder()
builder.setUid(uidValue)
builder.setToken(new SyncToken(5))
builder.setDeltaType(SyncDeltaType.CREATE)
builder.setObject(connectorObject)  // Use the ConnectorObjectBuilder class to build the ConnectorObject
 object.

// Call the handler with the SyncDelta object
handler builder.build()
```

- Using a Groovy Closure.

  In this case, the Closure delegates calls to a specific Object that implements these calls. For example:

```
handler {                       // The handler parameter here
                                                    is a Closure
    syncToken tokenValue        // (mandatory), the method resolution for 'syncToken' is delegated to
                                   the Object handling the Closure
    <DELTA_TYPE>()              // (mandatory), DELTA_TYPE should be one of: CREATE, UPDATE, DELETE,
                                   CREATE_OR_UPDATE
    object connectorObject      // (optional if DELTA_TYPE is a DELETE), the method resolution for
                                   'object' is delegated to the Object handling the Closure
    previousUid prevUidValue    // (optional), use only if UID has changed
}
```

In the following example, the handler is called twice - first for a CREATE and then for a DELETE:

```
// CREATE
handler({
    syncToken 15
    CREATE()
    object {
        id nameValue
        uid uidValue as String
        objectClass ObjectClass.GROUP
        attribute 'gid', gidValue
        attribute 'description', descriptionValue
    }
})

// DELETE
handler({
    syncToken 16
    DELETE(uidValue)
}
```

Optionally, when the action is SYNC, you might want to return a `SyncToken` at the end of the script. This is a convenient way to update the sync token if no relevant sync events are found.

6.4.9. Writing a Schema Script

A schema script builds the schema for the connector, either from a static, predefined schema, or by reading the schema from the external resource. The script should use the `builder` object to create the schema.

A sample schema script for an SQL database is provided in `openidm/samples/scripted-sql-with-mysql/tools/SchemaScript.groovy`

Input variables:

The following variables are available to a sync script:

configuration

> A handler to the connector's configuration object.

operation

> An OperationType that corresponds to the action (`SCHEMA`).

builder

> An instance of the `ICFObjectBuilder`. The `schema()` method should be called with a Closure parameter defining the schema objects.

> For more information, see "Using the `builder` Parameter".

log

>   A logger instance for the connector.

Returns: This script has no return value.

6.4.9.1. Using the `builder` Parameter

The `builder.schema()` must call the delegates `objectClass` method and `operationOptions` method to define the schema.

Call the `objectClass()` method for each object type (account, group, and so on) that must be defined. Include the call to the following delegates:

- `type()` - the name for this object type

- `attribute()` - define a single attribute for this object type

- `attributes()` - define multiple attribute for this object type

- `disable()` - list the operations for which this object type is forbidden

The following example defines a simple ACCOUNT object type:

```
builder.schema({
    objectClass {
        type ObjectClass.ACCOUNT_NAME
        attribute OperationalAttributeInfos.PASSWORD
        attribute PredefinedAttributeInfos.DESCRIPTION
        attribute 'groups', String.class, EnumSet.of(MULTIVALUED)
        attributes {
            userName String.class, REQUIRED
            email REQUIRED, MULTIVALUED
            __ENABLE__  Boolean.class
            createDate  NOT_CREATABLE, NOT_UPDATEABLE
            lastModified Long.class, NOT_CREATABLE, NOT_UPDATEABLE, NOT_RETURNED_BY_DEFAULT
            passwordHistory String.class, MULTIVALUED, NOT_UPDATEABLE, NOT_READABLE,
 NOT_RETURNED_BY_DEFAULT
            firstName()
            sn()
        }
    }
}
```

6.4.10. Writing a Resolve Username Script

A resolve username script resolves an object to its `uid` based on its `username`.

A sample resolve username script for an SQL database is provided in `openidm/samples/scripted-sql-with-mysql/tools/ResolveUsernameScript.groovy`

Input variables:

The following variables are available to a resolve username script:

configuration

A handler to the connector's configuration object.

options

A handler to the Operation Options.

operation

An OperationType that corresponds to the action (`RESOLVE_USERNAME`).

objectClass

The object class for which the username is resolved, such as `__ACCOUNT__` or `__GROUP__`.

username

A string that represents the username of the object.

log

A logger instance for the connector.

Returns: The user unique ID (OpenICF `__UID__`) of the object. The `type` of the returned UID must be a `string` or a `Uid`. If a null value or an object type other than `string` or `Uid` is returned, the script must throw an exception.

Resolve Username Script

```
def operation = operation as OperationType
def configuration = configuration as ScriptedConfiguration
def username = username as String
def log = log as Log
def objectClass = objectClass as ObjectClass
def options = options as OperationOptions
if (objectClass.is(ObjectClass.ACCOUNT_NAME)) {
    if (username.equals("TESTOK1")) {
        return new Uid("123")
    }
    throw new UnknownUidException();
}
```

6.4.11. Writing a Run On Resource Script

A run on resource script runs directly on the target resource that is managed by the connector.

A sample run on resource script for a connector that connects to DS over REST is provided in `openidm/samples/scripted-rest-with-dj/tools/ScriptOnResourceScript.groovy`

Input variables:

The following variables are available to a run on resource script:

configuration

A handler to the connector's configuration object.

options

A handler to the Operation Options.

operation

An OperationType that corresponds to the action (`RUNSCRIPTONRESOURCE`).

arguments

The `arguments (Map)` of the script (can be `null`).

log

A logger instance for the connector.

Returns: Any object that is returned by the script.

Run on Resource Script

```
import groovyx.net.http.RESTClient
import org.apache.http.client.HttpClient
import org.forgerock.openicf.connectors.scriptedrest.ScriptedRESTConfiguration
import org.forgerock.openicf.connectors.groovy.OperationType
import org.identityconnectors.common.logging.Log
import org.identityconnectors.framework.common.objects.OperationOptions

def operation = operation as OperationType
def configuration = configuration as ScriptedRESTConfiguration
def httpClient = connection as HttpClient
def connection = customizedConnection as RESTClient
def log = log as Log
def options = options as OperationOptions
def scriptArguments = scriptArguments as Map
def scriptLanguage = scriptLanguage as String
def scriptText = scriptText as String
```

6.4.12. Writing a Run On Connector Script

A run on connector script enables IDM to run a script in the context of the connector.

Input variables:

The following variables are available to a run on connector script:

configuration

> A handler to the connector's configuration object.

options

> A handler to the Operation Options.

operation

> An OperationType that corresponds to the action (`RUNSCRIPTONCONNECTOR`).

arguments

> The `arguments (Map)` of the script (can be `null`).

log

> A logger instance for the connector.

Returns: Any object that is returned by the script.

6.5. Advanced - Customizing the Configuration Initialization

Connectors created with the Groovy Connector Toolkit are, by default, stateful connectors. This means that the connector configuration instance is created only once.

The Groovy Connector Toolkit is precompiled code, and connector configurations are initialized in a specific way. If you have specific initialization requirements, you can customize the way in which the connector configuration instance is initialized, before the first script is evaluated. The `CustomizerScript.groovy` file enables you to define custom closures to interact with the default implementation.

The `CustomizerScript.groovy` file, provided with each compiled connector implementation, defines closures, such as `init {}`, `decorate {}`, and `destroy {}`. These closures are called during the lifecycle of the configuration.

When you unpack the Groovy Connector Toolkit JAR file, the `CustomizerScript.groovy` file is located at `org/forgerock/openicf/connectors/`*`connector-implementation`*.

FORGEROCK

Chapter 7
Writing Scripted Connectors With the PowerShell Connector Toolkit

You can use the PowerShell Connector Toolkit to create connectors that can provision any Microsoft system, including, but not limited to, Active Directory, MS SQL, MS Exchange, Sharepoint, Office365, and Azure. Essentially, any task that can be performed with PowerShell can be executed through connectors based on this toolkit.

The PowerShell Connector Toolkit is not a complete connector, in the traditional sense. Rather, it is a framework within which you must write your own PowerShell scripts to address the requirements of your Microsoft Windows ecosystem.

Connectors created with the PowerShell Connector Toolkit run on the .NET platform and require the installation of a .NET connector server on the Windows system. To install the .NET connector server, follow the instructions in "Installing and Configuring a .NET Connector Server" in the *Integrator's Guide*. These connectors also require PowerShell V2.

The PowerShell Connector Toolkit is available from ForgeRock Backstage. To install the connector, download the archive (`mspowershell-connector-1.4.1.0.zip`) and extract the `MsPowerShell.Connector.dll` to the same folder in which the Connector Server (`connectorserver.exe`) is located. IDM provides sample connector configurations and scripts that will enable you to get started with this toolkit.

7.1. About the PowerShell Scripting Language

PowerShell combines a command-line shell and scripting language, built on the .NET Framework. For a comprehensive introduction to PowerShell, and examples of how to use it, see Microsoft's TechNet article, Scripting with Windows PowerShell.

Chapter 8
Troubleshooting Connectors

Sometimes it is difficult to assess whether the root of a problem is at the OpenICF or connector level, or at the application level.

If you are using OpenICF connectors with IDM, you can adjust the log levels for specific parts of the system in the `path/to/openidm/conf/logging.properties` file.

The OpenICF API sets the `LoggingProxy` at a very high level. You can consider the Logging Proxy as the *border* between the application (OpenIDM) and the OpenICF framework.

To start a troubleshooting process, you should therefore enable the Logging Proxy and set it at a level high enough to provide the kind of information you need:

```
org.identityconnectors.framework.impl.api.LoggingProxy.level=FINE
```

```
org.identityconnectors.framework.impl.api.LoggingProxy.level=DEBUG
```

```
#Enable the LoggingProxy
org.identityconnectors.framework.impl.api.LoggingProxy.level=FINE

#Select the operation you want to trace, to trace all add:
org.identityconnectors.framework.api.operations.level=FINE

#To trace only some:
org.identityconnectors.framework.api.operations.CreateApiOp.level=FINE
org.identityconnectors.framework.api.operations.UpdateApiOp.level=FINE
org.identityconnectors.framework.api.operations.DeleteApiOp.level=FINE
```

The complete list of operations that you can trace is as follows:

```
AuthenticationApiOp
CreateApiOp
DeleteApiOp
GetApiOp
ResolveUsernameApiOp
SchemaApiOp
ScriptOnConnectorApiOp
ScriptOnResourceApiOp
SearchApiOp
SyncApiOp
TestApiOp
UpdateApiOp
ValidateApiOp
```

To enable logging in the remote Java Connector Server, edit the xml configuration file `/lib/framework/logback.xml` to uncomment the following line:

```
<logger name="org.identityconnectors.framework.impl.api.LoggingProxy" level="DEBUG" additivity="false">
    <appender-ref ref="TRACE-FILE"/>
</logger>
```

To enable logging in the remote .NET Connector Server, edit the configuration file `ConnectorServer.exe.config`, setting the following value to `true`

```
<add key="logging.proxy" value="false"/>
```

Index