



Samples Guide

ForgeRock Identity Management 5

Lana Frost
Mike Jang

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

Copyright © 2011-2017 ForgeRock AS.

Abstract

Guide providing a number of "sample deployments" that walk you through the essential features of ForgeRock® Identity Management software, as they would be implemented.



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

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

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

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

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

DejaVu Fonts

Bitstream Vera Fonts Copyright

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

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

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

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

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

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

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

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

Arev Fonts Copyright

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

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

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

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

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

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

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

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

Admonition graphics by Yannick Lung. Free for commercial use. Available at FreeCms.Cumulus.

Table of Contents

Preface	v
1. About This Guide	v
2. Formatting Conventions	v
3. Accessing Documentation Online	vi
4. Joining the ForgeRock Community	vii
1. Overview of the Samples	1
1.1. Samples Provided With IDM	1
1.2. Installing the Samples	6
1.3. Preparing the Server	6
2. Reconciling Data Between IDM and an XML File	7
2.1. Reconciling an XML File Resource	7
2.2. Using Scripts to Generate Log Messages	19
2.3. Demonstrating Asynchronous Reconciliation Using a Workflow	19
3. LDAP Samples - Reconciling Data Between IDM and LDAP Directories	23
3.1. Sample 2 - LDAP One Way	23
3.2. Sample 2b - LDAP Two Way	27
3.3. Sample 2c - Synchronizing LDAP Group Membership	32
3.4. Sample 2d - Synchronizing LDAP Groups	38
3.5. Sample 5 - Synchronization of Two LDAP Resources	42
3.6. Sample 5b - Failure Compensation With Multiple Resources	45
3.7. Sample 6 - LiveSync With an AD Server	48
3.8. Linking Historical Accounts	55
3.9. Storing Multiple Passwords For Managed Users	64
4. Samples That Use the Groovy Connector Toolkit to Create Scripted Connectors	78
4.1. Using the Connector Bundler to Build a ScriptedSQL Connector	78
4.2. Using the Groovy Connector Toolkit to Connect to OpenDJ With ScriptedREST	89
4.3. Using the Groovy Connector Toolkit to Connect to OpenDJ With ScriptedCREST	99
5. Samples That Use the PowerShell Connector Toolkit to Create Scripted Connectors	110
5.1. Connect to Active Directory	110
5.2. Connect to Azure AD	119
6. Audit Samples	138
6.1. Directing Audit Information To a MySQL Database	138
6.2. Show Audit Events Published on a JMS Topic	141
7. Scripted JMS Sample	145
7.1. Starting IDM and the ActiveMQ Broker	145
7.2. Access the REST Interface via the ActiveMQ UI	146
7.3. Customizing the Scripted JMS Sample	148
8. Demonstrating the Roles Implementation	151
8.1. CRUDOPS Role Sample - Working With Managed Roles	151
8.2. Provisioning Role Sample - Provisioning to an LDAP Server	162
8.3. Temporal Constraints Sample - Applying Time-Based Constraints to Roles ...	183

9. The Multi-Account Linking Sample	191
9.1. Before You Start: Link Qualifiers, Agents, and Insured Customers	191
9.2. External LDAP Configuration	193
9.3. Running the Multi-Account Linking Sample	193
9.4. Reconciling Managed Users to the External LDAP Server	202
10. The Trusted Servlet Filter Sample	203
10.1. Before You Start	203
10.2. The Sample Servlet Filter	203
10.3. Run the Sample	204
10.4. Create a Trusted User	204
10.5. Customizing the Sample for an External System	205
11. Integrating IDM With the ForgeRock Identity Platform	207
11.1. Preparing Your Systems	207
11.2. Preparing an Instance of OpenDJ	208
11.3. Starting IDM	208
11.4. Installing OpenAM for Integration	208
11.5. Configuring OpenAM for Integration with IDM	212
11.6. Plugging IDM Into OpenAM	218
11.7. Demonstrating Integration	221
11.8. The Integrated <code>authentication.json</code> File	222
11.9. Reconciliation and OpenAM	223
11.10. Integrating Social ID Logins	224
11.11. Registering Users via OpenIDM	225
12. Workflow Samples	227
12.1. Sample Workflow - Provisioning User Accounts	227
12.2. Workflow Use Cases	232
13. Google Sample - Connecting to Google With the Google Apps Connector	250
13.1. Before You Start	250
13.2. Configuring the Google Apps Connector	251
13.3. Running the Google Apps Sample	254
14. Connecting to Salesforce With the Salesforce Connector	261
14.1. Before you Start	261
14.2. Install the Sample	263
14.3. Running the Sample by Using the Admin UI	263
14.4. Running the Sample by Using the Command Line	268
15. Scripted Kerberos Connector Sample	274
15.1. Editing the Kerberos Connector Configuration	274
15.2. Running the Kerberos Sample	275
16. Customer Data Management Sample	283
16.1. Set Up Registration and Authentication	283
16.2. Setting Up Users for Marketo	284
16.3. Configuring the Marketo Connector	285
16.4. Reviewing Marketo Leads	286
17. Custom Endpoint Sample	287
Glossary	291
Index	294

Preface

ForgeRock Identity Platform™ is the only offering for access management, identity management, user-managed access, directory services, and an identity gateway, designed and built as a single, unified platform.

The platform includes the following components that extend what is available in open source projects to provide fully featured, enterprise-ready software:

- ForgeRock Access Management (AM)
- ForgeRock Identity Management (IDM)
- ForgeRock Directory Services (DS)
- ForgeRock Identity Gateway (IG)

1. About This Guide

This guide describes a number of sample deployments that demonstrate the core functionality of ForgeRock Identity Management software. The samples correspond to the configurations provided in the [openidm/samples](#) directory.

This guide is written for anyone testing ForgeRock Identity Management to manage identities, and to ensure compliance with identity management regulations.

The guide covers a number of ForgeRock Identity Management features, often including multiple features in a single sample.

You do not need a complete understanding of ForgeRock Identity Management software to learn something from this guide, although a background in identity management and maintaining web application software can help. You do need some background in managing services on your operating systems and in your application servers. You can nevertheless get started with this guide, and then learn more as you go along.

2. Formatting Conventions

Most examples in the documentation are created in GNU/Linux or Mac OS X operating environments. If distinctions are necessary between operating environments, examples are labeled with the operating environment name in parentheses. To avoid repetition file system directory names are

often given only in UNIX format as in `/path/to/server`, even if the text applies to `C:\path\to\server` as well.

Absolute path names usually begin with the placeholder `/path/to/`. This path might translate to `/opt/`, `C:\Program Files\`, or somewhere else on your system.

Command-line, terminal sessions are formatted as follows:

```
$ echo $JAVA_HOME
/path/to/jdk
```

Command output is sometimes formatted for narrower, more readable output even though formatting parameters are not shown in the command.

Program listings are formatted as follows:

```
class Test {
    public static void main(String [] args) {
        System.out.println("This is a program listing.");
    }
}
```

3. 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.
- ForgeRock core 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.

Core documentation therefore follows a three-phase review process designed to eliminate errors:

- Product managers and software architects review project documentation design with respect to the readers' software lifecycle needs.
- Subject matter experts review proposed documentation changes for technical accuracy and completeness with respect to the corresponding software.
- Quality experts validate implemented documentation changes for technical accuracy, completeness in scope, and usability for the readership.

The review process helps to ensure that documentation published for a ForgeRock release is technically accurate and complete.

Fully reviewed, published core documentation is available at <http://backstage.forgerock.com/>. Use this documentation when working with a ForgeRock Identity Platform release.

4. Joining the ForgeRock Community

Visit the [Community resource center](#) where you can find information about each project, download trial builds, browse the resource catalog, ask and answer questions on the forums, find community events near you, and find the source code for open source software.

Chapter 1

Overview of the Samples

This chapter lists all the samples provided with OpenIDM and gives a high-level overview of the purpose of each sample. This chapter also provides information that is required for all of the samples. Read this chapter, specifically Section 1.2, "Installing the Samples" and Section 1.3, "Preparing the Server" before you try any of the samples provided with OpenIDM.

1.1. Samples Provided With IDM

A number of samples are provided in the `openidm/samples` directory. This section describes the purpose of each sample:

Getting Started

The Getting Started sample describes how to install and evaluate OpenIDM.

Chapter 2, "Reconciling Data Between IDM and an XML File"

The XML samples all use the XML file connector to interact with an XML file resource. The samples demonstrate the following OpenIDM functionality:

- Section 2.1, "Reconciling an XML File Resource"

The basic XML sample demonstrates a connection to an XML file that holds user data. The sample includes one mapping and shows reconciliation from an XML file to the OpenIDM repository.

- Section 2.2, "Using Scripts to Generate Log Messages"

The logging sample demonstrates the logging capabilities available to OpenIDM scripts, and provides an alternative method for debugging scripts.

- Section 2.3, "Demonstrating Asynchronous Reconciliation Using a Workflow"

The workflow sample uses the XML connector to demonstrate asynchronous reconciliation using the workflow mechanism.

Chapter 3, "LDAP Samples - Reconciling Data Between IDM and LDAP Directories"

The LDAP samples all assume a connection to an LDAP directory, usually OpenDJ, or Active Directory. Samples 5 and 5b simulate an LDAP directory with an XML file, and use the XML connector. These samples demonstrate a wide variety of OpenIDM functionality and are broken down as follows:

- Section 3.1, "Sample 2 - LDAP One Way"

This sample uses the generic LDAP connector to connect to an LDAP directory. The sample includes one mapping from the LDAP directory to the managed user repository, and demonstrates reconciliation from the external resource to the repository.

- Section 3.2, "Sample 2b - LDAP Two Way"

This sample uses the generic LDAP connector to connect to an LDAP directory. The sample includes two mappings, one from the LDAP directory to the managed user repository, and one from the repository to the LDAP directory. The sample demonstrates reconciliation in both directions.

- Section 3.3, "Sample 2c - Synchronizing LDAP Group Membership"

This sample uses the generic LDAP connector to connect to an LDAP directory. The sample includes two mappings, one from the LDAP directory to the managed user repository, and one from the repository to the LDAP directory. The sample demonstrates synchronization of group membership, that is, how the value of the `ldapGroups` property in a managed user object is mapped to the corresponding user object in LDAP.

- Section 3.4, "Sample 2d - Synchronizing LDAP Groups"

This sample uses the generic LDAP connector to connect to an LDAP directory. The sample builds on the previous sample by providing an additional mapping, from the LDAP groups object, to the managed groups object. The sample illustrates a new managed object type (groups) and shows how this object type is synchronized with group containers in LDAP.

- Section 3.5, "Sample 5 - Synchronization of Two LDAP Resources"

Although this sample is grouped with the LDAP samples, it actually *simulates* two LDAP directories with XML files, and uses the XML file connector to connect the two. The purpose of this sample is to demonstrate reconciliation directly between two external resources, without the data passing through the OpenIDM repository. The sample also demonstrates the configuration of an outbound email service to send reconciliation summaries by mail.

- Section 3.6, "Sample 5b - Failure Compensation With Multiple Resources"

This sample builds on the previous sample to demonstrate a failure compensation mechanism that relies on script event hooks. The failure compensation mechanism ensures that reconciliation changes are propagated throughout a multiple-resource deployment, or rolled back in the case of error. The purpose of this mechanism is to keep the data consistent across multiple resources.

- Section 3.7, "Sample 6 - LiveSync With an AD Server"

This sample illustrates the liveSync mechanism that pushes changes from an external resource to the OpenIDM repository. The sample uses an LDAP connector to connect to an LDAP directory, either OpenDJ or Active Directory.

- Section 3.8, "Linking Historical Accounts"

This sample demonstrates the retention of inactive (historical) LDAP accounts that have been linked to a corresponding managed user account. The sample builds on sample 2b and uses the LDAP connector to connect to an OpenDJ instance. You can use any LDAP-v3 compliant directory server.

- Section 3.9, "Storing Multiple Passwords For Managed Users"

This sample demonstrates how to set up multiple passwords for managed users and how to synchronize separate passwords to different external resources. The sample includes two target LDAP servers, each with different password policy and encryption requirements. The sample also shows how to extend the password history policy to apply to multiple password fields.

Chapter 4, "Samples That Use the Groovy Connector Toolkit to Create Scripted Connectors"

The samples in this section use the Groovy Connector Toolkit to create a scripted connector. Because you can use scripted Groovy connectors to connect to a large variety of systems, the samples in this section show connections to several different external resources. The samples are broken down as follows:

- Section 4.1, "Using the Connector Bundler to Build a ScriptedSQL Connector"

This sample uses the *custom scripted connector bundler* to create a new scripted connector. The connector bundler generates a scripted connector, the connector configuration and the Groovy scripts required to communicate with an external MySQL database (HRDB).

- Section 4.2, "Using the Groovy Connector Toolkit to Connect to OpenDJ With ScriptedREST"

This sample uses the Groovy Connector Toolkit to implement a ScriptedREST connector, which interacts with the OpenDJ REST API.

- Section 4.3, "Using the Groovy Connector Toolkit to Connect to OpenDJ With ScriptedCREST"

This sample uses the Groovy Connector Toolkit to create a ScriptedCREST connector, which connects to an OpenDJ instance. The main difference between a CREST-based API and a generic REST API is that the CREST API is inherently recognizable by all ForgeRock products. As such, the sample can leverage CREST resources in the groovy scripts, to create CREST requests.

Chapter 5, "Samples That Use the PowerShell Connector Toolkit to Create Scripted Connectors"

The samples in this section use the PowerShell Connector Toolkit to create a scripted PowerShell connector. The samples are broken down as follows:

- Section 5.1, "Connect to Active Directory"

This sample uses the MS Active Directory PowerShell module to demonstrate how you can synchronize managed object data with a Microsoft Active Directory deployment. The sample provides a number of PowerShell scripts that enable you to perform basic CRUD (create, read, update, delete) operations on an Active Directory server.

- Section 5.2, "Connect to Azure AD"

This sample uses the Microsoft Azure Active Directory (Azure AD) PowerShell module to demonstrate how you can synchronize managed object data such as users and groups with a Microsoft AzureAD deployment.

Chapter 6, "Audit Samples"

The samples in this chapter demonstrate two ways to configure OpenIDM to save audit data:

- Section 6.1, "Directing Audit Information To a MySQL Database"

One audit sample uses a ScriptedSQL implementation of the Groovy Connector Toolkit to direct audit information to a MySQL database

- Section 6.2, "Show Audit Events Published on a JMS Topic"

The JMS audit sample demonstrates how the JMS audit event handler can publish messages that comply with the *Java(TM) Message Service Specification Final Release 1.1*

Chapter 7, "Scripted JMS Sample"

This sample demonstrates the OpenIDM scripted JMS message handler, and how it performs ForgeRock REST operations.

Chapter 8, "Demonstrating the Roles Implementation"

This sample builds on Section 3.1, "Sample 2 - LDAP One Way", and extends that sample to demonstrate how roles are implemented in OpenIDM.

Chapter 9, "The Multi-Account Linking Sample"

This sample illustrates how OpenIDM addresses links from multiple accounts to one identity.

Chapter 10, "The Trusted Servlet Filter Sample"

This sample demonstrates how to use a custom servlet filter and the Trusted Request Attribute Authentication Module in OpenIDM. Once configured, OpenIDM can use the servlet filter to authenticate through another service.

Chapter 11, "Integrating IDM With the ForgeRock Identity Platform"

This sample demonstrates the integration of three ForgeRock products: OpenIDM, OpenDJ, and OpenAM. With this sample, you can see how you can use OpenAM for authentication, for user identities that are maintained with OpenIDM, based on a data store of users in OpenDJ.

Chapter 12, "Workflow Samples"

The workflow sample and use cases demonstrate how OpenIDM uses workflows to provision user accounts. The samples demonstrate the use of the Self-Service UI to enable user self-registration,

- Section 12.1, "Sample Workflow - Provisioning User Accounts"

The provisioning workflow sample demonstrates a typical use case of a workflow — provisioning new users. The sample demonstrates the use of the Admin UI, to configure user self-service and the Self-Service UI that enables users to complete their registration process.

- Section 12.2, "Workflow Use Cases"

The workflow use cases work together to provide a complete business story, with the same set of sample data. Each of the use cases is integrated with the Self-Service UI.

Chapter 13, "Google Sample - Connecting to Google With the Google Apps Connector"

This sample uses the Google Apps Connector to manage the creation of users and groups on an external Google system, using OpenIDM's REST interface.

Chapter 14, "Connecting to Salesforce With the Salesforce Connector"

This sample uses the Salesforce Connector demonstrate reconciliation of user accounts from the OpenIDM repository to Salesforce, and from Salesforce to the OpenIDM repository.

Chapter 15, "Scripted Kerberos Connector Sample"

This sample demonstrates how to use the scripted Kerberos connector to manage Kerberos user principals and to reconcile user principals with OpenIDM managed user objects.

Chapter 16, "Customer Data Management Sample"

This sample demonstrates the CDM process to register and authenticate users with multiple social identity providers - Google, Facebook, and LinkedIn. The sample also demonstrates the use of the Marketo connector to reconcile this data to a Marketo database for lead generation (marketing).

Chapter 17, "Custom Endpoint Sample"

OpenIDM supports scriptable custom endpoints that enable you to launch arbitrary scripts through an OpenIDM REST URI. This sample shows how custom endpoints are configured and returns a list of variables available to each method used in a custom endpoint script.

In addition to these complete, runnable samples, the `samples/` directory includes the following subdirectories:

- `infoservice` - provides a sample custom health check script. This functionality is described in Section 2.3.4, "Customizing Health Check Scripts" in the *Integrator's Guide*.
- `misc` - includes sample configuration files for email (see Chapter 23, "Sending Email" in the *Integrator's Guide*), workflows (see Section 20.2, "Setting Up Activiti Integration" in the *Integrator's Guide*, and user self-service (see Chapter 5, "Configuring User Self-Service" in the *Integrator's Guide*).
- `provisioners` - includes sample connector configuration files.
- `schedules` - provides sample schedule configurations. This functionality is described in Section 16.2, "Configuring Schedules" in the *Integrator's Guide*.

- **security** - includes sample certificates and command-line examples for managing the keystore and truststore.
- **syncfailure** - includes configuration files and scripts to enable synchronization failure handling. This functionality is described in Section 13.3.5, "Setting the Synchronization Failure Configuration" in the *Integrator's Guide*.
- **taskscanner** - includes configuration files and scripts to demonstrate scheduled tasks on managed objects. This functionality is described in Section 16.7.1, "Configuring the Task Scanner" in the *Integrator's Guide*.

1.2. Installing the Samples

Each sample directory in `openidm/samples/` contains a number of subdirectories, such as `conf/` and `script/`. To start OpenIDM with a sample configuration, navigate to the `/path/to/openidm` directory and use the `-p` option of the **startup** command to point to the sample whose configuration you want to use. Some, but not all samples require additional software, such as an external LDAP server or database.

Many of the procedures in this guide refer to paths such as `samplex/...`. In each of these cases, the complete path is assumed to be `/path/to/openidm/samples/samplex/...`.

When you move from one sample to the next, bear in mind that you are changing the IDM configuration. For information on how configuration changes work, see Section 7.2, "Changing the Default Configuration" in the *Integrator's Guide*.

The command-line examples in this chapter (and throughout the OpenIDM documentation) assume a UNIX shell. If you are running these samples on Windows, adjust the command-line examples accordingly. For an indication of what the corresponding Windows command would look like, see the examples in Section 2.1, "Reconciling an XML File Resource".

1.3. Preparing the Server

Install an instance of IDM specifically to try the samples. That way you can experiment as much as you like, and discard the result if you are not satisfied.

If you are using the same instance for multiple samples, it is helpful to clear out the repository created for an earlier sample. To do so, shut down the server and delete the `openidm/db/openidm` directory.

```
$ rm -rf /path/to/openidm/db/openidm
```

IDM should then be ready to start with a new sample. For a number of the samples in this guide, users are created either with the UI or directly with a commons REST call. Users that have been created in the repository (managed users) should be able to log into the Self-Service UI.

Chapter 2

Reconciling Data Between IDM and an XML File

This chapter walks you through the XML samples (those samples labeled Sample 1, Sample 8, and Sample 9 in the `openidm/samples` directory). For a complete list of the samples provided with OpenIDM, and an overview of each sample, see Chapter 1, "*Overview of the Samples*".

2.1. Reconciling an XML File Resource

This chapter provides an overview of the first sample and how it is configured. For a complete list of the samples provided with OpenIDM, and an overview of each sample, see Chapter 1, "*Overview of the Samples*" or the README in the `openidm/samples` directory.

2.1.1. About the XML Sample

OpenIDM connects data objects held between resources by mapping one object to another. To connect to external resources, OpenIDM uses *connectors* that are configured for each external resource.

When objects in one external resource change, OpenIDM determines how the changes affect other objects, and can make the changes as necessary. This sample demonstrates how OpenIDM does this by using *reconciliation*. Reconciliation compares the objects in one resource to mapped objects in another resource. For a complete explanation of reconciliation and synchronization, see Section 14.1, "Types of Synchronization" in the *Integrator's Guide*.

In this sample, OpenIDM connects to an XML file that holds sample user data. The XML file is configured as the authoritative source. A *mapping* is configured between objects in the XML file and managed user objects in OpenIDM's repository.

Note that you can use OpenIDM to synchronize objects between two external resources without going through the OpenIDM repository. In such a case, objects are synchronized directly through connectors to the external resources.

This sample involves only one external resource. In practice, you can connect as many resources as needed for your deployment.

Sample Configuration Files

You can find configuration files for the sample under the `openidm/samples/sample1/conf` directory. As you review the sample, keep the following in mind:

- Start OpenIDM with the configuration associated with Sample 1:

```
$ ./startup.sh -p samples/sample1
```

For more information, see Section 2.1.2, "Install the Sample".

- OpenIDM regularly scans for any scheduler configuration files in the `conf` directory.
- OpenIDM's reconciliation service reads the mappings and actions for the source and target users from `conf/sync.json`.
- When you initiate a reconciliation, OpenIDM queries all users in the source, and then creates, deletes, or modifies users in the local OpenIDM repository as mapped in `conf/sync.json`.
- OpenIDM writes all operations to the audit logs in both the internal database and also the flat files in the `openidm/audit` directory.
- The default Sample 1 version of the `conf/authentication.json` file includes several authentication modules: `STATIC_USER`, `MANAGED_USER`, `INTERNAL_USER`, and `CLIENT_CERT`. For more information, see Section 18.1.2.2, "Supported Authentication Modules" in the *Integrator's Guide*.

When you start OpenIDM with the `-p` project variable (`./startup.sh -p samples/sample1`), the `&{launcher.project.location}` is set to a value of `samples/sample1`. The configuration files use this location, as shown in the following sections.

The following configuration files play important roles in this sample:

`samples/sample1/conf/provisioner.openicf-xml.json`

This file provides the configuration for this instance of the XML connector. It describes, among other things, the connector version, the location of the XML file resource, and the object types that are supported for this connection.

`samples/sample1/data/xmlConnectorData.xml`

This XML file is the external resource or data store in this sample. The XML file acts as the authoritative source for users. In the connector configuration file you can see that the `xmlFilePath` is set to `&{launcher.project.location}/data/xmlConnectorData.xml`.

The `&{launcher.project.location}`, in this case, is `samples/sample1`.

For details on connector configuration files, see Chapter 13, "Connecting to External Resources" in the *Integrator's Guide*.

`samples/sample1/conf/sync.json`

This file, also called a *mapping* file, defines the configuration for reconciliation and synchronization. This sample file includes only one mapping - `systemXmlAccounts_managedUser`. The mapping specifies the synchronization configuration between the XML file (source) and the OpenIDM repository (target). Examine the file to see how objects are mapped between the two resources, and the actions that OpenIDM should take when it finds objects in specific situations:

```
{
  "mappings": [
    {
      "name": "systemXmlfileAccounts_managedUser",
      "source": "system/xmlfile/account",
      "target": "managed/user",
      "correlationQuery": {
        "type": "text/javascript",
        "source": "var query = {'_queryId' : 'for-userName',
          'uid' : source.name};query;"
      },
      "properties": [
        {
          "source": "email",
          "target": "mail"
        },
        {
          "source": "firstname",
          "target": "givenName"
        },
        {
          "source": "lastname",
          "target": "sn"
        },
        {
          "source": "description",
          "target": "description"
        },
        {
          "source": "_id",
          "target": "_id"
        },
        {
          "source": "name",
          "target": "userName"
        },
        {
          "source": "password",
          "target": "password"
        },
        {
          "source": "mobileTelephoneNumber",
          "target": "telephoneNumber"
        },
        {
          "source": "roles",
          "transform": {
            "type": "text/javascript",
            "source": "var _ = require('lib/lodash'); _.map(source.split(','),
              function(role) { return {'_ref': 'repo/internal/role/' + role} });"
          }
        }
      ]
    }
  ]
}
```



```
    },
    "target" : "authzRoles"
  }
],
"policies": [
  {
    "situation": "CONFIRMED",
    "action": "UPDATE"
  },
  {
    "situation": "FOUND",
    "action": "IGNORE"
  },
  {
    "situation": "ABSENT",
    "action": "CREATE"
  },
  {
    "situation": "AMBIGUOUS",
    "action": "IGNORE"
  },
  {
    "situation": "MISSING",
    "action": "IGNORE"
  },
  {
    "situation": "SOURCE_MISSING",
    "action": "IGNORE"
  },
  {
    "situation": "UNQUALIFIED",
    "action": "IGNORE"
  },
  {
    "situation": "UNASSIGNED",
    "action": "IGNORE"
  }
]
}
]
```

Source and target paths that start with `managed`, such as `managed/user`, always refer to objects in the OpenIDM repository. Paths that start with `system`, such as `system/xmlfile/account`, refer to external objects, in this case, objects in the XML file.

For more information about synchronization, reconciliation, and `sync.json`, see Chapter 14, "Synchronizing Data Between Resources" in the *Integrator's Guide*.

`samples/sample1/conf/schedule-reconcile_systemXmlAccounts_managedUser.json`

The sample schedule configuration file defines a task that launches a reconciliation every minute for the mapping named `systemXmlAccounts_managedUser`. The schedule is disabled by default:

```
{
  "enabled" : false,
  "type": "cron",
  "schedule": "0 0/1 * * * ?",
  "persisted" : true,
  "misfirePolicy" : "fireAndProceed",
  "invokeService": "sync",
  "invokeContext": {
    "action": "reconcile",
    "mapping": "systemXmlfileAccounts_managedUser"
  }
}
```

OpenIDM regularly scans the `conf/` directory for any schedule configuration files. For information about the schedule configuration, see Chapter 16, "Scheduling Tasks and Events" in the *Integrator's Guide*.

Apart from the scheduled reconciliation run, you can also start the reconciliation run through the REST interface. The call to the REST interface is an HTTP POST such as the following:

```
$ curl \
  --header "X-OpenIDM-Username: openidm-admin" \
  --header "X-OpenIDM-Password: openidm-admin" \
  --request POST \
  "http://localhost:8080/openidm/recon?
_action=recon&mapping=systemXmlfileAccounts_managedUser&waitForCompletion=true"
```

The `waitForCompletion=true` parameter specifies that the operation should return only when it has completed.

2.1.2. Install the Sample

Start OpenIDM with the configuration for Sample 1:

```
$ cd /path/to/openidm
$ ./startup.sh -p samples/sample1
```

2.1.3. Review the Sample in the Administrative User Interface

OpenIDM includes a web-based Administrative User Interface, known as the Admin UI. For details, see Section 4.1, "Configuring the Server from the Admin UI" in the *Integrator's Guide*.

After starting OpenIDM, you can access the Admin UI by navigating to <https://localhost:8443/admin>. The first time you log in, use the default administrative credentials, (Login: `openidm-admin`, Password: `openidm-admin`).

Warning

To protect your deployment in production, change the default administrative password. To do so, navigate to the Self-Service UI at <https://localhost:8443/> and click Change Password.

You should now see the Dashboard screen, with quick start cards for common administrative tasks with the connectors and managed objects associated with that configuration.

2.1.4. Running Reconciliation

Reconcile the objects in the resources, either by setting `"enabled" : true` in the schedule configuration file (`conf/schedule-reconcile_systemXmlAccounts_managedUser.json`) and then waiting until the scheduled reconciliation happens, or by using the REST interface, as shown in the following example:

```
$ curl \
  --header "X-OpenIDM-Username: openidm-admin" \
  --header "X-OpenIDM-Password: openidm-admin" \
  --request POST \
  "http://localhost:8080/openidm/recon?
  _action=recon&mapping=systemXmlfileAccounts_managedUser&waitForCompletion=true"
```

Successful reconciliation returns a reconciliation run ID, and the status of the reconciliation operation, as follows:

```
{
  "_id": "2d87c817-3d00-4776-a705-7de2c65937d8",
  "state": "SUCCESS"
}
```

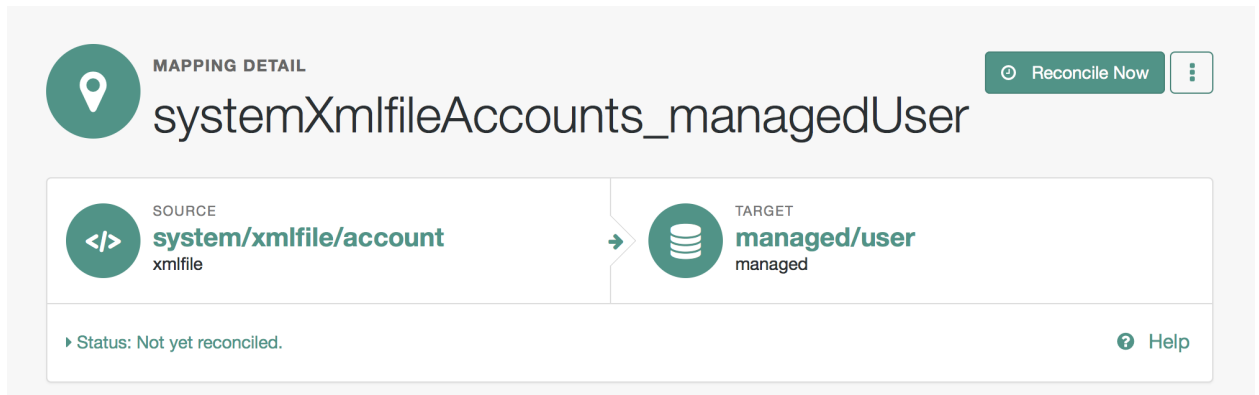
Alternatively, you can run the same reconciliation in the Admin UI:

1. Click Configure > Mappings.

For Sample 1, you should see one mapping, `systemXmlfileAccounts_managedUser`.

2. Select Edit to access the configuration options associated with reconciliation.
3. To run the reconciliation, click Reconcile Now.

Figure 2.1. OpenIDM Admin UI Mappings with Sample 1



2.1.5. Viewing Users and Logs

After reconciliation, you can use the Admin UI to display user records in both the source and target resources:

1. Navigate to the URL where OpenIDM is installed.

If it is local, navigate to <https://localhost:8443/admin>.

2. Click Configure > Mappings, then select the only available mapping ([systemXmlfileAccounts_managedUser](#))
3. On the Association tab, you should see the result of the reconciliation, from source to target, at the bottom of the screen.

You can also use the REST interface to display all users in the local repository. Use a REST client to perform an HTTP GET on the following URL: http://localhost:8080/openidm/managed/user?_queryId=query-all-ids with the headers "X-OpenIDM-Username: openidm-admin" and "X-OpenIDM-Password: openidm-admin".

OpenIDM returns JSON data. Depending on the browser, you can use a REST client to display the JSON or download it as a file. Alternatively, you can use the following **curl** command to get the JSON response:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/managed/user?_queryId=query-all-ids"

{
  "result": [
    {
      "_id": "scarter",
      "_rev": "1"
    },
    {
      "_id": "bjensen",
      "_rev": "1"
    }
  ]
},
...
}
```

In addition to querying the users by their ID, you can set up arbitrary queries. For more information about using query expressions in a REST call, see Section 8.3, "Defining and Calling Queries" in the *Integrator's Guide*.

Now try a RESTful GET of user `bjensen` by appending the user ID to the managed user URL (<http://localhost:8080/openidm/managed/user/>):

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/managed/user/bjensen"
{
  "_id": "bjensen",
  "_rev": "1",
  "mail": "bjensen@example.com",
  "givenName": "Barbara",
  "sn": "Jensen",
  "description": "Created By XML1",
  "userName": "bjensen@example.com",
  "telephoneNumber": "1234567",
  "accountStatus": "active",
  "effectiveRoles": [],
  "effectiveAssignments": []
}
```

The complete user record is returned. If you need this level of information for all users, substitute `query-all` for `query-all-ids`.

You can filter the output with the query expressions described in Section 8.3, "Defining and Calling Queries" in the *Integrator's Guide*.

As defined in the mapping file `conf/sync.json`, the `sn` and `mail` parameters correspond to surname (`sn`) and email address, respectively.

For example, the following RESTful GET filters output by surname (sn):

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/managed/user?_queryFilter=true&_fields=sn"

{
  "result": [
    {
      "_id": "scarter",
      "_rev": "1",
      "sn": "Carter"
    },
    {
      "_id": "bjensen",
      "_rev": "1",
      "sn": "Jensen"
    }
  ]
},
...
}
```

Now that you have a list of users, you can add more fields to your query:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/managed/user?_queryFilter=true&_fields=sn,mail,description"

{
  "result": [
    {
      "_id": "scarter",
      "_rev": "1",
      "sn": "Carter",
      "mail": "scarter@example.com",
      "description": "Created By XML1"
    },
    {
      "_id": "bjensen",
      "_rev": "1",
      "sn": "Jensen",
      "mail": "bjensen@example.com",
      "description": "Created By XML1"
    }
  ]
},
...
}
```

This information is also available in the CSV format audit logs located in the `openidm/audit` directory:

```
$ ls /path/to/openidm/audit/
access.csv activity.csv recon.csv
```

For more information about the contents of each of these files, see Section 21.3, "Audit Log Event Topics" in the *Integrator's Guide*.

You can get a similar level of information for each user. For example, after running reconciliation, follow the instructions in Section 2.1.5, "Viewing Users and Logs", and review information from the reconciled linked resource.

2.1.6. Adding Users in a Resource

Add a user to the source connector XML data file to see reconciliation in action. During the next reconciliation, OpenIDM finds the new user in the source connector, and creates the user in the local repository.

1. To add the user copy the following XML into `openidm/samples/sample1/data/xmlConnectorData.xml`:

```
<ri: __ACCOUNT__ >
  <icf: __UID__ >tmorris</icf: __UID__ >
  <icf: __NAME__ >tmorris@example.com</icf: __NAME__ >
  <ri:firstname>Toni</ri:firstname>
  <ri:lastname>Morris</ri:lastname>
  <ri:email>tmorris@example.com</ri:email>
  <ri:mobileTelephoneNumber>1234567</ri:mobileTelephoneNumber>
  <ri:roles>openidm-authorized</ri:roles>
  <icf: __DESCRIPTION__ >Created By XML</icf: __DESCRIPTION__ >
</ri: __ACCOUNT__ >
```

2. Run reconciliation again, as described in Section 2.1.4, "Running Reconciliation".
3. After reconciliation has run, query the local repository to see the new user appear in the list of all managed users:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/managed/user?_queryId=query-all-ids"
{
  "result": [
    {
      "_id": "scarter",
      "_rev": "2"
    },
    {
      "_id": "bjensen",
      "_rev": "2"
    },
    {
      "_id": "tmorris",
      "_rev": "1"
    }
  ]
  ,
  ...
}
```

To see what happened during the reconciliation operation, look at the reconciliation audit log, `openidm/audit/recon.csv`. This formatted excerpt from the log covers the two reconciliation runs done in this sample:

```
"_id", "action", ..., "reconId", "situation", "sourceObjectId", "targetObjectId", "timestamp";
"7e...", "CREATE", ..., "486...", "ABSENT", "system/xmlfile/acc.../bjensen", "managed/user/bjensen", ...;
"1a...", "CREATE", ..., "486...", "ABSENT", "system/xmlfile/acc.../scarter", "managed/user/scarter", ...;
"33...", "UPDATE", ..., "aa9...", "CONFIRMED", "system/xmlfile/acc.../bjensen", "managed/user/bjensen", ...;
"1d...", "UPDATE", ..., "aa9...", "CONFIRMED", "system/xmlfile/acc.../scarter", "managed/user/scarter", ...;
"0e...", "CREATE", ..., "aa9...", "ABSENT", "system/xmlfile/acc.../tmorris", "managed/user/tmorris", ...;
```

The relevant audit log fields in this example are: `action`, `situation`, `sourceObjectId`, and `targetObjectId`. For each object in the source, reconciliation leads to an action on the target.

In the first reconciliation run (abbreviated `reconID` is shown as `486...`), the source object does not exist in the target, resulting in an `ABSENT` situation and an action to `CREATE` the object in the target. The object created earlier in the target does not exist in the source, and so is `IGNORED`.

In the second reconciliation run (abbreviated `reconID` is shown as `aa9...`), after you added a user to the source XML, OpenIDM performs an `UPDATE` on the user objects `bjensen` and `scarter` that already exist in the target. OpenIDM performs a `CREATE` on the target for the new user (`tmorris`).

You configure the action that OpenIDM takes based on an object's situation in the configuration file, `conf/sync.json`. For the list of all possible situations and actions, see Chapter 14, "*Synchronizing Data Between Resources*" in the *Integrator's Guide*.

For details about auditing, see Chapter 21, "*Logging Audit Information*" in the *Integrator's Guide*.

2.1.7. Adding Users Over REST

You can add users to the local repository over the REST interface. The following example adds a user named James Berg.

Create `james` (UNIX):


```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--header "Content-Type: application/json" \
--request POST \
--data '{
  "_id": "jberg",
  "userName": "jberg",
  "sn": "Berg",
  "givenName": "James",
  "mail": "jberg@example.com",
  "telephoneNumber": "5556787",
  "description": "Created by OpenIDM REST.",
  "password": "MyPassw0rd"
}' \
"http://localhost:8080/openidm/managed/user?_action=create"
{
  "_id": "jberg",
  "_rev": "1",
  "userName": "jberg",
  "sn": "Berg",
  "givenName": "James",
  "mail": "jberg@example.com",
  "telephoneNumber": "5556787",
  "description": "Created by OpenIDM REST.",
  "accountStatus": "active",
  "effectiveRoles": [],
  "effectiveAssignments": []
}
```

Create `james` (Windows):

```
C:\> curl ^
--header "X-OpenIDM-Username: openidm-admin" ^
--header "X-OpenIDM-Password: openidm-admin" ^
--header "Content-Type: application/json" ^
--request POST ^
--data "{ \"_id\": \"jberg\", \"userName\": \"jberg\", \"sn\": \"Berg\", \"givenName\": \"James\", \"email\": \"jberg@example.com\", \"telephoneNumber\": \"5556787\", \"description\": \"Created by OpenIDM REST.\", \"password\": \"MyPassw0rd\" }" ^
"http://localhost:8080/openidm/managed/user?_action=create"
```

The output is essentially the same as the UNIX command output.

OpenIDM creates the new user in the repository. If you configure a mapping to apply changes from the local repository to the XML file connector as a target, OpenIDM then updates the XML file to add the new user.

You can also add users through the UI, which uses the OpenIDM REST API. When you have logged into the UI as the OpenIDM administrator, click Manage > User > New User. The process is straightforward.

2.2. Using Scripts to Generate Log Messages

OpenIDM provides a `logger` object with `debug()`, `error()`, `info()`, `trace()`, and `warn()` functions that you can use to log messages to the OSGi console from your scripts.

2.2.1. Install the Sample

Prepare OpenIDM as described in Section 1.3, "Preparing the Server", then start OpenIDM with the configuration for sample 8.

```
$ cd /path/to/openidm
$ ./startup.sh -p samples/sample8
```

The `sync.json` file in the `sample8/conf` directory includes brief examples of log messages.

2.2.2. Running the Sample

Run reconciliation over the REST interface.

```
$ curl \
  --header "X-OpenIDM-Username: openidm-admin" \
  --header "X-OpenIDM-Password: openidm-admin" \
  --request POST \
  "http://localhost:8080/openidm/recon?
  _action=recon&mapping=systemXmlfileAccounts_managedUser&waitForCompletion=true"
```

The reconciliation operation returns a reconciliation run ID, and the status of the operation.

Note the log messages displayed in the OSGi console. The following example omits timestamps and so forth to show only the message strings.

```
->
...Case no Source: the source object contains: = null [5235432-...
...Case emptySource: the source object contains: = {lastname=Carter, mobile...
...Case sourceDescription: the source object contains: = Created By XML1
...Case onCreate: the source object contains: = {lastname=Carter, mobile...
...Case result: the source object contains: = {SOURCE_IGNORED={count=0, ids=[]},...
```

2.3. Demonstrating Asynchronous Reconciliation Using a Workflow

Sample 9 demonstrates asynchronous reconciliation using workflows. Reconciliation generates an approval request for each ABSENT user. The configuration for this action is defined in the `conf/sync.json` file, which specifies that an `ABSENT` condition should launch the `managedUserApproval` workflow:

```
...
{
  "situation" : "ABSENT",
  "action" : {
    "workflowName" : "managedUserApproval",
    "type" : "text/javascript",
    "file" : "workflow/triggerWorkflowFromSync.js"
  },
  ...
}
```

When the request is approved by an administrator, the absent users are created by an asynchronous reconciliation process.

Prepare a fresh installation of OpenIDM before trying this sample.

2.3.1. Install the Sample

Prepare OpenIDM as described in Section 1.3, "Preparing the Server", then start OpenIDM with the configuration for sample 9.

```
$ cd /path/to/openidm
$ ./startup.sh -p samples/sample9
```

2.3.2. Running the Sample

1. Run reconciliation over the REST interface.

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request POST \
"http://localhost:8080/openidm/recon?
_action=recon&mapping=systemXmlfileAccounts_managedUser&waitForCompletion=true"
```

The reconciliation operation returns a reconciliation run ID, and the status of the operation.

Reconciliation starts an approval workflow for each ABSENT user. These approval workflows (named `managedUserApproval`) wait for the request to be approved by an administrator.

2. Query the invoked workflow task instances over REST.

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/workflow/taskinstance?_queryId=query-all-ids"
```

In this case, the request returns two workflow results, each with a process ID (`_id`) as well as a process definition ID. You will use the value of the `_id` shortly.

```
{
  "result" : [ {
```

```

"tenantId" : "",
"createTime" : "2014-05-01T13:48:42.980-08:00",
"executionId" : "101",
"delegationStateString" : null,
"processVariables" : { },
  "_id" : "123",
"processInstanceId" : "101",
"description" : null,
"priority" : 50,
"name" : "Evaluate request",
"dueDate" : null,
"parentTaskId" : null,
"processDefinitionId" : "managedUserApproval:1:3",
"taskLocalVariables" : { },
"suspensionState" : 1,
"assignee" : "openidm-admin",
"cachedElContext" : null,
"queryVariables" : null,
"activityInstanceVariables" : { },
"deleted" : false,
"suspended" : false,
  "_rev" : 1,
"revisionNext" : 2,
"category" : null,
"taskDefinitionKey" : "evaluateRequest",
"owner" : null,
"eventName" : null,
"delegationState" : null
}, {
"tenantId" : "",
"createTime" : "2014-05-01T13:48:42.980-08:00",
"executionId" : "102",
"delegationStateString" : null,
"processVariables" : { },
  "_id" : "124",
"processInstanceId" : "102",
"description" : null,
"priority" : 50,
"name" : "Evaluate request",
"dueDate" : null,
"parentTaskId" : null,
"processDefinitionId" : "managedUserApproval:1:3",
"taskLocalVariables" : { },
"suspensionState" : 1,
"assignee" : "openidm-admin",
"cachedElContext" : null,
"queryVariables" : null,
"activityInstanceVariables" : { },
"deleted" : false,
"suspended" : false,
  "_rev" : 1,
"revisionNext" : 2,
"category" : null,
"taskDefinitionKey" : "evaluateRequest",
"owner" : null,
"eventName" : null,
"delegationState" : null
} ],
"resultCount" : 2,

```

```
"pagedResultsCookie" : null,  
"remainingPagedResults" : -1  
}
```

3. Approve the requests over REST, by setting the `requestApproved` parameter for the specified task instance to `true`. Note the use of one of the values of `_id` in the REST call, in this case, `124`.

On UNIX:

```
$ curl \  
--header "X-OpenIDM-Username: openidm-admin" \  
--header "X-OpenIDM-Password: openidm-admin" \  
--header "Content-Type: application/json" \  
--request POST \  
--data '{"requestApproved": "true"}' \  
"http://localhost:8080/openidm/workflow/taskinstance/124?_action=complete"
```

On Windows:

```
$ curl ^  
--header "X-OpenIDM-Username: openidm-admin" ^  
--header "X-OpenIDM-Password: openidm-admin" ^  
--header "Content-Type: application/json" ^  
--request POST ^  
--data "{\"requestApproved\": \"true\"}" ^  
"http://localhost:8080/openidm/workflow/taskinstance/124?_action=complete"
```

A successful call returns the following:

```
{"Task action performed":"complete"}
```

4. Once the request has been approved, an asynchronous reconciliation operation runs, which creates the users whose accounts were approved in the previous step.

List the users that were created by the asynchronous reconciliation.

```
$ curl \  
--header "X-OpenIDM-Username: openidm-admin" \  
--header "X-OpenIDM-Password: openidm-admin" \  
--request GET \  
"http://localhost:8080/openidm/managed/user?_queryId=query-all-ids"
```

One user is returned.

```
{  
  "result": [ {  
    "_rev": "0",  
    "_id": "1"  
  } ],  
  ...  
}
```

Chapter 3

LDAP Samples - Reconciling Data Between IDM and LDAP Directories

This chapter walks you through the LDAP samples (those samples labeled 2, 2b, 2c, 2d, 5, 5b and 6 in the `openidm/samples` directory). For a complete list of the samples provided with OpenIDM, and an overview of each sample, see Chapter 1, "*Overview of the Samples*".

3.1. Sample 2 - LDAP One Way

Sample 2 resembles the Section 2.1, "Reconciling an XML File Resource", but in sample 2 OpenIDM is connected to a local LDAP server. The sample has been tested with OpenDJ, but should work with any LDAPv3-compliant server.

Sample 2 demonstrates how OpenIDM can pick up new or changed objects from an external resource. The sample contains only one mapping, from the external LDAP server resource to the OpenIDM repository. The sample therefore does not push any changes made to OpenIDM managed user objects out to the LDAP server.

3.1.1. LDAP Server Configuration

Sample 2 expects the following configuration for the external LDAP server:

- The LDAP server runs on the local host.
- The LDAP server listens on port 1389.
- A user with DN `cn=Directory Manager` and password `password` has read access to the LDAP server.
- Directory data for that server is stored under base DN `dc=example,dc=com`.
- User objects for that server are stored under base DN `ou=People,dc=example,dc=com`.
- User objects have the object class `inetOrgPerson`.
- User objects have the following attributes:
 - `cn`
 - `description`

- `givenName`
- `mail`
- `sn`
- `telephoneNumber`
- `uid`
- `userPassword`

An example user object follows.

```
dn: uid=jdoe,ou=People,dc=example,dc=com
objectClass: person
objectClass: organizationalPerson
objectClass: inetOrgPerson
objectClass: top
givenName: John
uid: jdoe
cn: John Doe
telephoneNumber: 1-415-523-0772
sn: Doe
mail: jdoe@example.com
description: Created by OpenIDM
userPassword: password
```

The following steps provide setup instructions for OpenDJ directory server 5. Adjust these instructions if you are using an older version of OpenDJ, or an alternative LDAP server.

1. Download OpenDJ from ForgeRock's BackStage site and extract the zip archive.

The LDIF data for this sample is provided in the file `openidm/samples/sample2/data/Example.ldif`. Import this data during your OpenDJ setup.

2. Set up OpenDJ and import the `Example.ldif` file for this sample:

```

$ cd /path/to/openssl
$ ./setup \
  directory-server \
  --rootUserDN "cn=Directory Manager" \
  --rootUserPassword password \
  --hostname localhost \
  --ldapPort 1389 \
  --adminConnectorPort 4444 \
  --baseDN dc=com \
  --ldifFile /path/to/openidm/samples/sample2/data/Example.ldif \
  --acceptLicense
Validating parameters .....Done.
Configuring Certificates .....Done.
Configuring server .....Done.
Importing data from /path/to/openidm/samples/sample2/data/Example.ldif .....Done.
Starting Directory Server .....Done.

To see basic server status and configuration, you can launch /path/to/openssl/bin/status

```

3.1.2. Install the Sample

Prepare OpenIDM as described in Section 1.3, "Preparing the Server", then start OpenIDM with the configuration for sample 2.

```

$ cd /path/to/openidm
$ ./startup.sh -p samples/sample2

```

3.1.3. Reconcile the Repository

The mapping configuration file ([sync.json](#)) for this sample includes the mapping `systemLdapAccounts_managedUser`, which synchronizes users from the source LDAP server with the target OpenIDM repository.

You can run this part of the sample by using the `curl` command-line utility, or by using the OpenIDM Administration UI. This section provides instructions for both methods.

Procedure 3.1. Run the Sample Using the Command Line

1. Reconcile the repository by running the following command:

```

$ curl \
  --header "X-OpenIDM-Username: openidm-admin" \
  --header "X-OpenIDM-Password: openidm-admin" \
  --request POST \
  "http://localhost:8080/openidm/recon?
  _action=recon&mapping=systemLdapAccounts_managedUser&waitForCompletion=true"
{
  "_id": "b1394d10-29b0-4ccf-81d8-c88948ea121c-4",
  "state": "SUCCESS"
}

```


The reconciliation operation creates the two users from the LDAP server in the OpenIDM repository, assigning the new objects random unique IDs.

2. To retrieve the users from the repository, query their IDs as follows:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/managed/user?_queryId=query-all-ids"
{
  "result": [
    {
      "_id": "f52df646-7108-45e1-9342-1a17f257b497",
      "_rev": "1"
    },
    {
      "_id": "f7fccf54-e76a-404c-93f0-7486d30f1dc3",
      "_rev": "1"
    }
  ]
  ,
  ...
}
```

3. To retrieve individual user objects, include the ID in the URL, for example:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/managed/user/0a5546d6-149b-4f8b-b3be-4afa8a267d45"
{
  "_id": "f7fccf54-e76a-404c-93f0-7486d30f1dc3",
  "_rev": "1",
  "displayName": "Barbara Jensen",
  "description": "Created for OpenIDM",
  "givenName": "Barbara",
  "mail": "bjensen@example.com",
  "sn": "Jensen",
  "telephoneNumber": "1-360-229-7105",
  "userName": "bjensen",
  "accountStatus": "active",
  "effectiveRoles": [],
  "effectiveAssignments": []
}
```

Procedure 3.2. Run the Sample Using the Admin UI

1. Log in to the Admin UI at the URL <https://localhost:8443/admin> as the default administrative user (`openidm-admin`) with password `openidm-admin`.

Warning



To protect your deployment in production, change the default administrative password. To do so, navigate to the Self-Service UI at <https://localhost:8443/> and click Change Password.

2. Click Configure > Mappings.

This page shows one configured mapping, from the `ldap` server to the OpenIDM repository (`managed/user`).

Mappings Help

[+ New Mapping](#) Filter...

SOURCE  System/Ldap/Account ldap	→	TARGET  Managed/User managed		
systemLdapAccounts_managedUser Not yet reconciled.	Move	Edit	Delete	Reconcile

3. Click Reconcile.

The reconciliation operation creates the two users from the LDAP server in the OpenIDM repository.

4. Retrieve the users in the repository. Click Manage > User.

You should now see two users from the LDAP server, reconciled to the OpenIDM repository.

5. When you click a username, you can view the details of that user account.

3.2. Sample 2b - LDAP Two Way

Like sample 2, sample 2b connects to an external LDAP server. However, sample 2b has two mappings configured, one from the LDAP server to the OpenIDM repository, and the other from the OpenIDM repository to the LDAP server.

3.2.1. External LDAP Configuration

As demonstrated for sample 2, you can use OpenDJ as an LDAP server. The LDIF data for this sample is provided in the file `openidm/samples/sample2b/data/Example.ldif`. Import this data during your OpenDJ setup.

Configure the LDAP server as for sample 2, Section 3.1.1, "LDAP Server Configuration", but import the LDIF file that is specific to Sample 2b during the setup. The LDAP user must have write access to create users from OpenIDM on the LDAP server.

3.2.2. Install the Sample

Prepare OpenIDM as described in Section 1.3, "Preparing the Server", then start OpenIDM with the configuration for sample 2b.

```
$ cd /path/to/openidm
$ ./startup.sh -p samples/sample2b
```

3.2.3. Run the Sample

The mapping configuration file (`sync.json`) for this sample includes two mappings, `systemLdapAccounts_managedUser`, which synchronizes users from the source LDAP server with the target OpenIDM repository, and `managedUser_systemLdapAccounts`, which synchronizes changes from the OpenIDM repository to the LDAP server.

You can run this part of the sample by using the `curl` command-line utility, or by using the OpenIDM Administration UI. This section provides instructions for both methods.

Procedure 3.3. Run the Sample Using the Command Line

1. Reconcile the repository over the REST interface by running the following command:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request POST \
"http://localhost:8080/openidm/recon?
_action=recon&mapping=systemLdapAccounts_managedUser&waitForCompletion=true"
{
  "state": "SUCCESS",
  "_id": "027e25e3-7a33-4858-9080-161c2b40a6bf-2"
}
```

The reconciliation operation returns a reconciliation run ID and the status of the operation. Reconciliation creates user objects from LDAP in the OpenIDM repository, assigning the new objects random unique IDs.

2. To retrieve the users from the repository, query their IDs as follows:

```

$ curl \
  --header "X-OpenIDM-Username: openidm-admin" \
  --header "X-OpenIDM-Password: openidm-admin" \
  --request GET \
  "http://localhost:8080/openidm/managed/user?_queryId=query-all-ids"
{
  "result": [
    {
      "_id": "d460ed00-74f9-48fb-8cc1-7829be60ddac",
      "_rev": "1"
    },
    {
      "_id": "74fe2d25-4eb1-4148-a3ae-ff80f194b3a6",
      "_rev": "1"
    }
  ]
},
...
}

```

3. To retrieve individual user objects, include the ID in the URL, for example:

```

$ curl \
  --header "X-OpenIDM-Username: openidm-admin" \
  --header "X-OpenIDM-Password: openidm-admin" \
  --request GET \
  "http://localhost:8080/openidm/managed/user/d460ed00-74f9-48fb-8cc1-7829be60ddac"
{
  "_id": "d460ed00-74f9-48fb-8cc1-7829be60ddac",
  "_rev": "1",
  "displayName": "Barbara Jensen",
  "description": "Created for OpenIDM",
  "givenName": "Barbara",
  "mail": "bjensen@example.com",
  "telephoneNumber": "1-360-229-7105",
  "sn": "Jensen",
  "userName": "bjensen",
  "accountStatus": "active",
  "effectiveRoles": [],
  "effectiveAssignments": []
}

```

4. Test the second mapping by creating a user in the OpenIDM repository.

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--header "Content-Type: application/json" \
--request POST \
--data '{
  "mail":"fdoe@example.com",
  "sn":"Doe",
  "telephoneNumber":"555-1234",
  "userName":"fdoe",
  "givenName":"Felicitas",
  "description":"Felicitas Doe",
  "displayName":"fdoe"}' \
"http://localhost:8080/openidm/managed/user?_action=create"
{
  "_id": "90d1f388-d8c3-4438-893c-be4e498e7a1c",
  "_rev": "1",
  "mail": "fdoe@example.com",
  "sn": "Doe",
  "telephoneNumber": "555-1234",
  "userName": "fdoe",
  "givenName": "Felicitas",
  "description": "Felicitas Doe",
  "displayName": "fdoe",
  "accountStatus": "active",
  "effectiveRoles": [],
  "effectiveAssignments": []
}
```

5. By default, *implicit synchronization* is enabled for mappings from the `managed/user` repository to any external resource. This means that when you update a managed object, any mappings defined in the `sync.json` file that have the managed object as the source are automatically executed to update the target system. For more information, see Section 14.3.2, "Mapping Source Objects to Target Objects" in the *Integrator's Guide*.

Test that the implicit synchronization has been successful by querying the users in the LDAP directory over REST, as follows:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/system/ldap/account?_queryId=query-all-ids"
{
  "result": [
    {
      "_id": "uid=jdoe,ou=People,dc=example,dc=com",
      "dn": "uid=jdoe,ou=People,dc=example,dc=com"
    },
    {
      "_id": "uid=bjensen,ou=People,dc=example,dc=com",
      "dn": "uid=bjensen,ou=People,dc=example,dc=com"
    },
    {
      "_id": "uid=fdoe,ou=People,dc=example,dc=com",
      "dn": "uid=fdoe,ou=People,dc=example,dc=com"
    }
  ]
  ,
  ...
}
```

Note the new entry for user `fdoe`.

- Query the complete entry by including `fdoe`'s ID in the URL.

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/system/ldap/account/uid=fdoe,ou=People,dc=example,dc=com"
{
  "_id": "uid=fdoe,ou=People,dc=example,dc=com",
  "mail": "fdoe@example.com",
  "employeeType": null,
  "ldapGroups": [],
  "telephoneNumber": "555-1234",
  "givenName": "Felicitas",
  "cn": "fdoe",
  "dn": "uid=fdoe,ou=People,dc=example,dc=com",
  "uid": "fdoe",
  "sn": "Doe",
  "description": "Felicitas Doe"
}
```

Procedure 3.4. Run the Sample Using the Admin UI

- Log in to the Admin UI at the URL <https://localhost:8443/admin> as the default administrative user (`openidm-admin`) with password `openidm-admin`.

Warning

To protect your deployment in production, change the default administrative password. To do so, navigate to the Self-Service UI at <https://localhost:8443/> and click Change Password.

2. Click Configure > Mappings.

This tab shows two configured mappings, one from the `ldap` server to the OpenIDM repository (`managed/user`) and one from the OpenIDM repository to the `ldap` server.

3. On the first mapping (LDAP to managed user), click Reconcile.

The reconciliation operation creates the two users from the LDAP server in the OpenIDM repository.

4. Retrieve the users in the repository. Click Manage > User.
5. You should see two users from the LDAP server, reconciled to the OpenIDM repository.
6. To retrieve the details of a specific user, click that username in the User List page.
7. Add a new user in the OpenIDM repository by clicking New User in the User List page.

Complete the user details and click Save.

8. By default, *implicit synchronization* is enabled for mappings from the `managed/user` repository to any external resource. This means that when you update a managed object, any mappings defined in the `sync.json` file that have the managed object as the source are automatically executed to update the target system. For more information, see Section 14.3.2, "Mapping Source Objects to Target Objects" in the *Integrator's Guide*.

To test that the implicit synchronization has been successful, look at `fdoe`'s record, and click the Linked Systems tab. The information under this tab includes the external resource to which this user entry is mapped.

3.3. Sample 2c - Synchronizing LDAP Group Membership

Like sample 2b, sample 2c connects to an external LDAP server and has mappings from the LDAP server to the OpenIDM repository, and from the OpenIDM repository to the LDAP server. However, in sample 2c, LDAP group memberships are synchronized, in addition to user entries.

As demonstrated for sample 2, you can use OpenDJ as an LDAP server. The LDIF data for this sample is provided in the file `openidm/samples/sample2c/data/Example.ldif`.

3.3.1. External LDAP Configuration

Configure the LDAP server as for sample 2, Section 3.1.1, "LDAP Server Configuration". The LDAP user must have write access to create users from OpenIDM on the LDAP server. When you configure the LDAP server, import the LDIF file customized for this sample, `openidm/samples/sample2c/data/Example.ldif`. This file includes a number of LDAP groups, including the following:

```
dn: ou=Groups,dc=example,dc=com
ou: Groups
objectClass: organizationalUnit
objectClass: top

dn: cn=openidm,ou=Groups,dc=example,dc=com
uniqueMember: uid=jdoe,ou=People,dc=example,dc=com
cn: openidm
objectClass: groupOfUniqueNames
objectClass: top

dn: cn=openidm2,ou=Groups,dc=example,dc=com
uniqueMember: uid=bjensen,ou=People,dc=example,dc=com
cn: openidm2
objectClass: groupOfUniqueNames
objectClass: top
```

The users with DNS `uid=jdoe,ou=People,dc=example,dc=com` and `uid=bjensen,ou=People,dc=example,dc=com` are also imported with the `Example.ldif` file.

3.3.2. Install the Sample

Prepare OpenIDM as described in Section 1.3, "Preparing the Server", then start OpenIDM with the configuration for sample 2c.

```
$ cd /path/to/openidm
$ ./startup.sh -p samples/sample2c
```

3.3.3. Run the Sample

The mapping configuration file (`sync.json`) for this sample includes two mappings, `systemLdapAccounts_managedUser`, which synchronizes users from the source LDAP server with the target OpenIDM repository, and `managedUser_systemLdapAccounts`, which synchronizes changes from the OpenIDM repository to the LDAP server.

You can run this part of the sample by using the `curl` command-line utility, or by using the OpenIDM Administration UI. This section provides instructions for both methods.

Procedure 3.5. Run the Sample Using the Command Line

1. Reconcile the repository over the REST interface by running the following command:


```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request POST \
"http://localhost:8080/openidm/recon?
action=recon&mapping=systemLdapAccounts_managedUser&waitForCompletion=true"
{
  "_id": "6652c292-5309-40e5-b272-b74d67dd95c9-4",
  "state": "SUCCESS"
}
```

The reconciliation operation returns a reconciliation run ID and the status of the operation. Reconciliation creates user objects from LDAP in the OpenIDM repository, assigning the new objects random unique IDs.

2. To retrieve the users from the repository, query their IDs as follows:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/managed/user?_queryId=query-all-ids"
{
  "result": [
    {
      "_id": "b63fb9a7-99bc-4eb4-8bfd-15f14a756e5b",
      "_rev": "1"
    },
    {
      "_id": "8462fe0c-2ab2-459a-a25e-474474889c9e",
      "_rev": "1"
    }
  ]
  ,
  ...
}
```

3. To retrieve individual user objects, include the ID in the URL. The following request retrieves the user object for John Doe:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/managed/user/8462fe0c-2ab2-459a-a25e-474474889c9e"
{
  "_id": "8462fe0c-2ab2-459a-a25e-474474889c9e",
  "_rev": "1",
  "displayName": "John Doe",
  "description": "Created for OpenIDM",
  "givenName": "John",
  "mail": "jdoe@example.com",
  "telephoneNumber": "1-415-599-1100",
  "sn": "Doe",
  "userName": "jdoe",
  "ldapGroups": [
    "cn=openidm,ou=Groups,dc=example,dc=com"
  ],
  "accountStatus": "active",
  "effectiveRoles": [],
  "effectiveAssignments": []
}
```

Note that John Doe's user object contains an `ldapGroups` property, the value of which indicates his groups on the LDAP server:

```
"ldapGroups": ["cn=openidm,ou=Groups,dc=example,dc=com"]
```

4. Update John Doe's `ldapGroups` property, to change his membership from the `openidm` group to the `openidm2` group.

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--header "Content-Type: application/json" \
--request POST \
--data '[
  {
    "operation": "replace",
    "field": "/ldapGroups",
    "value": ["cn=openidm2,ou=Groups,dc=example,dc=com"]
  }
]' \
"http://localhost:8080/openidm/managed/user?_action=patch&_queryId=for-userName&uid=jdoe"
{
  "displayName": "John Doe",
  "description": "Created for OpenIDM",
  "givenName": "John",
  "mail": "jdoe@example.com",
  "telephoneNumber": "1-415-599-1100",
  "sn": "Doe",
  "userName": "jdoe",
  "accountStatus": "active",
  "effectiveRoles": [],
  "effectiveAssignments": [],
  "_id": "8462fe0c-2ab2-459a-a25e-474474889c9e",
  "_rev": "2",
  "ldapGroups": [
    "cn=openidm2,ou=Groups,dc=example,dc=com"
  ]
}
```

This command changes John Doe's `ldapGroups` property in the OpenIDM repository, from `"cn=openidm,ou=Groups,dc=example,dc=com"` to `"cn=openidm2,ou=Groups,dc=example,dc=com"`. As a result of implicit synchronization, the change is propagated to the LDAP server. John Doe is removed from the first LDAP group and added to the second LDAP group in OpenDJ.

5. You can verify this change by querying John Doe's record on the LDAP server, as follows:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/system/ldap/account/uid=jdoe,ou=People,dc=example,dc=com"
{
  "_id": "uid=jdoe,ou=People,dc=example,dc=com",
  "telephoneNumber": "1-415-599-1100",
  "description": "Created for OpenIDM",
  "sn": "Doe",
  "dn": "uid=jdoe,ou=People,dc=example,dc=com",
  "ldapGroups": [
    "cn=openidm2,ou=Groups,dc=example,dc=com"
  ],
  "uid": "jdoe",
  "cn": "John Doe",
  "givenName": "John",
  "mail": "jdoe@example.com"
}
```

Procedure 3.6. Run the Sample Using the Admin UI

1. Log in to the Admin UI at the URL <https://localhost:8443/admin> as the default administrative user (`openidm-admin`) with password `openidm-admin`.

Warning

To protect your deployment in production, change the default administrative password. To do so, navigate to the Self-Service UI at <https://localhost:8443/> and click Change Password.

2. Click Configure > Mappings.

This window shows two configured mappings, one from the `ldap` server to the OpenIDM repository (`managed/user`) and one from the OpenIDM repository to the `ldap` server.

3. On the first mapping (LDAP to managed user), click Reconcile.

The reconciliation operation creates the two users from the LDAP server in the OpenIDM repository.

4. Click Manage > User. Examine the users reconciled from the LDAP server to the internal repository.
5. Examine the two users from the LDAP server that have been reconciled to the OpenIDM repository.
6. To retrieve the details of a specific user, click that username. In this case, click on user `jdoe`.

Examine the information stored for user `john.doe`. Click the Linked Systems tab. The Linked Resource item indicates the external resource on which John Doe's managed object is mapped, in this case, `ldap account`.

In this linked resource, John Doe's `ldapGroups` are displayed. Currently, John Doe is a member of `cn=openidm,ou=Groups,dc=example,dc=com`.

7. Update John Doe's `ldapGroups` property to change his membership from the `openidm` group to the `openidm2` group. Currently, you can only do this over the REST interface, as follows:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--header "Content-Type: application/json" \
--request POST \
--data '[
  {
    "operation": "replace",
    "field": "/ldapGroups",
    "value": ["cn=openidm2,ou=Groups,dc=example,dc=com"]
  }
]' \
"http://localhost:8080/openidm/managed/user?_action=patch&_queryId=for-userName&uid=jdoe"
{
  "displayName": "John Doe",
  "description": "Created for OpenIDM",
  "givenName": "John",
  "mail": "jdoe@example.com",
  "telephoneNumber": "1-415-599-1100",
  "sn": "Doe",
  "userName": "jdoe",
  "accountStatus": "active",
  "effectiveRoles": [],
  "effectiveAssignments": [],
  "_id": "8462fe0c-2ab2-459a-a25e-474474889c9e",
  "_rev": "2",
  "ldapGroups": [
    "cn=openidm2,ou=Groups,dc=example,dc=com"
  ]
}
```

This command changes John Doe's `ldapGroups` property in the OpenIDM repository, from `cn=openidm,ou=Groups,dc=example,dc=com` to `cn=openidm2,ou=Groups,dc=example,dc=com`. As a result of implicit synchronization, the change is propagated to the LDAP server. John Doe is removed from the first LDAP group and added to the second LDAP group in OpenDJ.

8. You can verify this change by reloading John Doe's user information, clicking Linked Systems, and examining the value of his `ldapGroups` property.

3.4. Sample 2d - Synchronizing LDAP Groups

Sample 2d also connects to an external LDAP server and focuses on synchronization of LDAP groups.

Like sample 2, this sample uses OpenDJ directory server 5 but it should work with any LDAP v3-compliant LDAP server.

3.4.1. External LDAP Configuration

Configure the LDAP server as for sample 2, Section 3.1.1, "LDAP Server Configuration". The LDAP user must have write access to create users from OpenIDM on the LDAP server.

When you set up OpenDJ, import the LDIF data for this sample, from [openidm/samples/sample2d/data/Example.ldif](#). The import creates a number of LDAP groups, including the following:

```
dn: ou=Groups,dc=example,dc=com
ou: Groups
objectClass: organizationalUnit
objectClass: top

dn: cn=openidm,ou=Groups,dc=example,dc=com
uniqueMember: uid=jdoe,ou=People,dc=example,dc=com
cn: openidm
objectClass: groupOfUniqueNames
objectClass: top

dn: cn=openidm2,ou=Groups,dc=example,dc=com
uniqueMember: uid=bjensen,ou=People,dc=example,dc=com
cn: openidm2
objectClass: groupOfUniqueNames
objectClass: top
```

The user with dn `uid=jdoe,ou=People,dc=example,dc=com` is also imported with the `Example.ldif` file.

There is an additional user, `bjensen` in the sample LDIF file. This user is essentially a "dummy" user, provided for compliance with RFC 4519, which stipulates that every `groupOfUniqueNames` object must contain at least one `uniqueMember`. `bjensen` is not actually used in this sample.

3.4.2. Install the Sample

Prepare OpenIDM as described in Section 1.3, "Preparing the Server", then start OpenIDM with the configuration for sample 2d.

```
$ cd /path/to/openidm
$ ./startup.sh -p samples/sample2d
```

3.4.3. Running the Sample

The mapping configuration file (`sync.json`) for this sample includes three mappings:

- `systemLdapAccounts_managedUser`

Synchronizes users from the source LDAP server with the target OpenIDM repository,

- `managedUser_systemLdapAccounts`

Synchronizes changes from the OpenIDM repository to the LDAP server.

- `systemLdapGroups_managedGroup`

Synchronizes groups from the source LDAP server with the target OpenIDM repository.

Due to the similarity with other OpenIDM samples, especially samples 2b and 2c, the focus of this sample is on the mapping unique to this sample, `systemLdapGroups_managedGroup`.

You can run this part of the sample by using the `curl` command-line utility, or by using the OpenIDM Administration UI. This section provides instructions for both methods.

Procedure 3.7. Run the Sample Using the Command Line

1. Reconcile the repository over the REST interface for the group mapping, `systemLdapGroups_managedGroup` by running the following command:

```
$ curl \
  --header "X-OpenIDM-Username: openidm-admin" \
  --header "X-OpenIDM-Password: openidm-admin" \
  --request POST \
  "http://localhost:8080/openidm/recon?
  _action=recon&mapping=systemLdapGroups_managedGroup&waitForCompletion=true"
```

The reconciliation operation returns a reconciliation run ID, and the status of the operation.

2. The reconciliation creates managed group objects for each group that exists in OpenDJ. To list the managed groups, run the following command:

```
$ curl \
  --header "X-OpenIDM-Username: openidm-admin" \
  --header "X-OpenIDM-Password: openidm-admin" \
  --request GET \
  "http://localhost:8080/openidm/managed/group?_queryFilter=true"
```

The resulting JSON object should include content similar to the following.

```
{
  "result": [
    {
      "_id": "b6c4d7ce-2103-42c2-b5f2-74ca9309ad37",
      "_rev": "1",
      "dn": "cn=Contractors,ou=Groups,dc=example,dc=com",
      "description": null,
      "uniqueMember": [],
      "name": "Contractors"
    },
    {
      "_id": "2326b9ee-6975-4c19-aa3c-d228afc4ff71",
      "_rev": "1",
      "dn": "cn=openidm2,ou=Groups,dc=example,dc=com",
      "description": null,
      "uniqueMember": [
        "uid=bjensen,ou=People,dc=example,dc=com"
      ]
    }
  ]
}
```

```

    },
    "name": "openidm2"
  },
  {
    "_id": "035f6444-bce3-4931-96b7-e10b2301fe74",
    "_rev": "1",
    "dn": "cn=Employees,ou=Groups,dc=example,dc=com",
    "description": null,
    "uniqueMember": [],
    "name": "Employees"
  },
  {
    "_id": "65c8fb86-01e6-4fca-9237-e50c251f4575",
    "_rev": "1",
    "dn": "cn=Chat Users,ou=Groups,dc=example,dc=com",
    "description": null,
    "uniqueMember": [],
    "name": "Chat Users"
  },
  {
    "_id": "5c3e4965-16d7-4a8f-af73-3ab165b66cf9",
    "_rev": "1",
    "dn": "cn=openidm,ou=Groups,dc=example,dc=com",
    "description": null,
    "uniqueMember": [
      "uid=jdoe,ou=People,dc=example,dc=com"
    ],
    "name": "openidm"
  }
],
...
}

```

Procedure 3.8. Run the Sample Using the Admin UI

1. Log in to the Admin UI at the URL <https://localhost:8443/admin> as the default administrative user (`openidm-admin`) with password `openidm-admin`.

Warning

To protect your deployment in production, change the default administrative password. To do so, navigate to the Self-Service UI at <https://localhost:8443/> and click Change Password.

2. Click Configure > Mappings.

This page shows three configured mappings, from the `ldap` server accounts repository to the OpenIDM repository (`managed/user`), from the OpenIDM repository back to the `ldap` server, and from the `ldap` server group accounts repository to the OpenIDM `managed/group` repository.

3. On the third mapping (LDAP groups to managed groups), click Reconcile.

The reconciliation operation creates the two groups from the LDAP server in the OpenIDM repository.

- Retrieve the groups in the repository by clicking the Association tab below the mapping. Scroll down to Data Association Management.

▾ **Data Association Management**

Data displayed may be affected by recent/ongoing reconciliation or synchronization. Please run reconciliation to see the latest information

View: ALL SITUATIONS

Change Source to Target Association

Reconcile Selected Record

SOURCE	LINK QUALIFIER	TARGET
Filter... cn=Chat Users,ou=Groups,dc=example,dc=com Chat Users	default	Filter... cn=Chat Users,ou=Groups,dc=example,dc=com Chat Users
cn=Contractors,ou=Groups,dc=example,dc=com Contractors	default	cn=Contractors,ou=Groups,dc=example,dc=com Contractors
cn=openidm,ou=Groups,dc=example,dc=com uid=jdoe,ou=People,dc=example,dc=com openidm	default	cn=openidm,ou=Groups,dc=example,dc=com uid=jdoe,ou=People,dc=example,dc=com openidm
cn=openidm2,ou=Groups,dc=example,dc=com uid=bjensen,ou=People,dc=example,dc=com openidm2	default	cn=openidm2,ou=Groups,dc=example,dc=com uid=bjensen,ou=People,dc=example,dc=com openidm2
cn=Employees,ou=Groups,dc=example,dc=com Employees	default	cn=Employees,ou=Groups,dc=example,dc=com Employees

The five groups from the LDAP server (source) have been reconciled to the OpenIDM repository (target).

3.5. Sample 5 - Synchronization of Two LDAP Resources

Sample 5 demonstrates the flow of data from one external resource to another. The resources are named **LDAP** and **AD** but in the sample, both resources are simulated with XML files.

You can optionally configure an outbound email service, if you want to receive emailed reconciliation summaries, as described in the following section.

3.5.1. Configure Email for the Sample

If you do not configure the email service, the functionality of the sample does not change. However, you might see the following message in the OSGi console when you run a reconciliation operation:

```
Email service not configured; report not generated.
```

To configure OpenIDM to send a reconciliation summary by email, follow these steps:

1. Copy the template `external.email.json` file from the `samples/misc` directory to the `conf` directory of Sample 5:

```
$ cd /path/to/openidm
$ cp samples/misc/external.email.json samples/sample5/conf
```

2. Edit the `external.email.json` file for outbound email, as described in Chapter 23, "Sending Email" in the *Integrator's Guide*.
3. Edit the `reconStats.js` file from the `sample5/script` directory. Near the start of the file, configure the OpenIDM email service to send statistics to the email addresses of your choice:

```
var email = {
  //UPDATE THESE VALUES
  from : "openidm@example.com",
  to : "youremail@example.com",
  cc : "idmadmin2@example.com,idmadmin3@example.com",
  subject : "Recon stats for " + global.mappingName,
  type : "text/html"
},
template,
...
```

3.5.2. Install the Sample

No external configuration is required for this sample. Prepare OpenIDM as described in Section 1.3, "Preparing the Server", then start OpenIDM with the configuration of sample 5.

```
$ cd /path/to/openidm
$ ./startup.sh -p samples/sample5
```

The XML files that simulate the resources are located in the `openidm/samples/sample5/data/` folder. When you start OpenIDM with the sample 5 configuration, OpenIDM creates the `xml_AD_Data.xml` file, which does not contain users until you run reconciliation.

3.5.3. Run the Sample

Run a reconciliation operation, to synchronize the contents of the simulated LDAP resource to the OpenIDM repository.

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request POST \
"http://localhost:8080/openidm/recon?
_action=recon&mapping=systemLdapAccounts_managedUser&waitForCompletion=true"
```

This command creates a user in the repository. It is not necessary to run a second reconciliation operation to synchronize the AD resource. Automatic synchronization propagates any change made to managed users in the OpenIDM repository to the simulated AD resource.

Review the contents of `xml_AD_Data.xml`. It should now contain information for the same user that was present in the startup version of the `xml_LDAP_Data.xml` file.

Alternatively, you can list users in the AD resource with the following command:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/system/ad/account?_queryId=query-all-ids"

{
  "result" : [ {
    "name" : "DD0E1",
    "__UID__" : "8dad9df3-820d-41ea-a3ab-a80c241bbc98",
    "_id" : "8dad9df3-820d-41ea-a3ab-a80c241bbc98"
  } ]
,
...
}
```

You can use the `_id` of the user to read the user information from the AD resource, for example:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/system/ad/account/8dad9df3-820d-41ea-a3ab-a80c241bbc98"
{
  "email" : [ "mail1@example.com" ],
  "name" : "DD0E1",
  "__UID__" : "8dad9df3-820d-41ea-a3ab-a80c241bbc98",
  "firstname" : "Darth",
  "lastname" : "Doe",
  "_id" : "8dad9df3-820d-41ea-a3ab-a80c241bbc98"
}
```

To verify that the sample is working, repeat the process. Set up a second user in the `xml_LDAP_Data.xml` file. An example of how that file might appear with a second user (`GD0E1`) is shown here:

```
<?xml version="1.0" encoding="UTF-8"?>
<icf:OpenICFContainer
  xmlns:icf="http://openidm.forgerock.com/xml/ns/public/resource/openicf/resource-schema-1.xsd"
  xmlns:ri="http://openidm.forgerock.com/xml/ns/public/resource/instances/resource-schema-extension"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://openidm.forgerock.com/xml/ns/public/resource/instances/resource-schema-
extension
  /path/to/openidm/samples/sample5/data/resource-schema-extension.xsd
  http://openidm.forgerock.com/xml/ns/public/resource/openicf/resource-schema-1.xsd
  /path/to/openidm/samples/sample5/data/resource-schema-1.xsd">
  <ri:__ACCOUNT__>
    <icf:__UID__>1</icf:__UID__>
    <icf:__PASSWORD__>TestPassw0rd2</icf:__PASSWORD__>
    <ri:firstname>Darth</ri:firstname>
    <icf:__DESCRIPTION__>Created By XML1</icf:__DESCRIPTION__>
    <icf:__NAME__>DDOE1</icf:__NAME__>
    <ri:email>mail1@example.com</ri:email>
    <ri:lastname>Doe</ri:lastname>
  </ri:__ACCOUNT__>
  <ri:__ACCOUNT__>
    <icf:__UID__>2</icf:__UID__>
    <icf:__PASSWORD__>TestPassw0rd2</icf:__PASSWORD__>
    <ri:firstname>Garth</ri:firstname>
    <icf:__DESCRIPTION__>Created By XML1</icf:__DESCRIPTION__>
    <icf:__NAME__>GDOE1</icf:__NAME__>
    <ri:email>mail2@example.com</ri:email>
    <ri:lastname>Doe</ri:lastname>
  </ri:__ACCOUNT__>
</icf:OpenICFContainer>
```

Rerun the reconciliation and query REST commands shown previously. The reconciliation operation creates the new user from the simulated LDAP resource in the OpenIDM repository. An implicit synchronization operation then creates that user in the AD resource.

3.6. Sample 5b - Failure Compensation With Multiple Resources

The compensated synchronization mechanism depicted in this sample can help manage the risks associated with synchronizing data across multiple resources.

Typically, when a managed/user object is changed, implicit synchronization replays that change to all configured external resources. If synchronization fails for one target resource (for example, due to a policy validation failure on the target, or the target being unavailable), the synchronization operation stops at that point. The effect is that a record might be changed in the repository, and in the targets on which synchronization was successful, but not on the failed target, or any targets that would have been synchronized after the failure. This situation can result in disparate data sets across resources. While a reconciliation operation would eventually bring all targets back in sync, reconciliation can be an expensive operation with large data sets.

The compensated synchronization mechanism ensures that either all resources are synchronized successfully, or that the original change is rolled back. This mechanism uses an `onSync` script hook

configured with a `compensate.js` script that can be used to "revert" the partial change to managed/user and to the corresponding external resources.

Sample 5b is similar to sample 5 in that it simulates two external resources with XML files (located in the `sample5b/data` directory). The `xml_LDAP_Data.xml` file simulates an LDAP data source. OpenIDM creates the `xml_AD_Data.xml` file when you start OpenIDM with the sample. Sample 5b adds the `onSync` script hook to the process, configured in the `sample5b/conf/managed.json` file.

The following excerpt of the `managed.json` file shows the `onSync` hook, which calls the `compensate.js` script, provided in the `/path/to/openidm/bin/defaults/script` directory.

```
...
},
"onSync" : {
  "type" : "text/javascript",
  "file" : "compensate.js"
},
```

You can use the `onSync` script hook to ensure that changes made in the repository are synchronized to all external resources, or that no changes are made. For more information about how implicit synchronization uses the `onSync` script hook, see Section 14.12, "Configuring Synchronization Failure Compensation" in the *Integrator's Guide*.

You can optionally configure an outbound email service for this sample, if you want to receive emailed reconciliation summaries. The email service configuration is identical to that of Section 3.5.1, "Configure Email for the Sample".

3.6.1. Install the Sample

No external configuration is required for this sample. Prepare OpenIDM as described in Section 1.3, "Preparing the Server", then start OpenIDM with the configuration of sample 5b.

```
$ cd /path/to/openidm
$ ./startup.sh -p samples/sample5b
```

The XML files that simulate an external LDAP and AD resource are now located in the `openidm/samples/sample5b/data/` directory. The simulated AD data store file, `xml_AD_Data.xml`, does not contain users until you run reconciliation.

Run the sample in exactly the same way that you did for Sample 5, following the steps in Section 3.5.3, "Run the Sample". Those steps will reconcile a user to your internal managed user repository.

Unless you run the steps in Section 3.5.3, "Run the Sample", you will not be able to run the steps in the next section.

3.6.2. Demonstrate onSync

To demonstrate integration of the samples with the OpenIDM UI, this sample uses the UI to view and make changes to user objects in the repository. You can also use the REST interface to make these changes.

Log into the OpenIDM UI as the administrative user. On a local system, navigate to <https://localhost:8443/admin>. The default administrative account and password are both `openidm-admin`.

Select Manage > User. Make a change to the data of an existing user (`DDOE1`). As a result of the implicit synchronization from the managed object repository, that change is reflected almost immediately on the external resources. For sample 5b, you should see the changes in both XML files in the `sample5b/data` directory. Alternatively, you can query the external resources over REST, as described previously.

The synchronization is successful, across all configured external resources, so the updated user record can be seen in both the `xml_LDAP_Data.xml` and `xml_AD_Data.xml` files.

The next step is to simulate a problem connecting to the LDAP resource. One way to do so on the local system is to rename the LDAP data file so that it is unreadable. On a Linux system, the following command, as an administrative user, would serve that purpose:

```
$ cd /path/to/openidm/samples/sample5b/data
$ sudo mv xml_LDAP_Data.xml xml_LDAP_Data.xml.bak
```

In the UI, now try another update to user `DDOE1`. With the modified filename of the simulated LDAP resource, implicit synchronization cannot write to this resource. An error similar to the following is displayed in the log file, `openidm0.log.0`:

```
Data file does not exist:
/path/to/openidm/samples/sample5b/data/xml_LDAP_Data.xml
```

Although the AD resource is available, implicit synchronization will not reach this resource, because the mapping is specified *after* the managed/user to LDAP mapping in the `sync.json` file.

When the implicit synchronization operation fails for the LDAP resource, the `onSync` hook invokes the `compensate.js` script. This script attempts to revert the original change by performing another update to `DDOE1` in the repository (managed/user). This change, in turn, triggers another automatic synchronization to the AD and LDAP resources.

Because the LDAP resource is still unreadable, the synchronization to LDAP fails again, which triggers the `compensate.js` script again. This time, however, the script recognizes that the change was originally called as a result of a compensation and aborts.

The original synchronization error from the first update is thrown from the script and the UI should display that error. If you refresh the UI, and view that user entry again, you will notice that the change to the entry has been reverted.

Note that if you change the name of the AD resource file (to make it unavailable), a change to a managed/user entry will be synchronized successfully with the LDAP resource (because that mapping

appears first in `sync.json`). The synchronization will fail for the AD resource. In this case, the change will be reverted on both the managed/user entry, and the LDAP resource.

3.7. Sample 6 - LiveSync With an AD Server

Sample 6 resembles sample 5, but demonstrates liveSync from an external resource. Sample 6 includes configuration files for two scenarios, depending on whether you have a live Active Directory (AD) service, or whether you need to simulate an AD service with an OpenDJ server. Each scenario is associated with a file in the `sample6/alternatives` directory. Depending on your scenario, copy the corresponding file to the `sample6/conf` directory.

Active AD Deployment

If you have an actual AD deployment available, copy the `provisioner.openicf-realad.json` file to the `conf/` subdirectory. You can then configure synchronization between an OpenDJ Directory Server and an active AD deployment.

As this sample demonstrates synchronization *from* the AD server *to* OpenDJ, data on the AD server is not changed.

Simulated AD Deployment

If you need to simulate an AD deployment, copy the `provisioner.openicf-fakead.json` file to the `conf/` subdirectory. You can then configure synchronization between an OpenDJ Directory server and a simulated AD server.

This sample simulates an AD server on the same instance of OpenDJ, using a different base DN.

The options shown in the associated configuration files can be easily modified to work with any standard LDAP server.

3.7.1. Sample 6 with an Active AD Deployment

If you have an existing, active instance of AD, set up OpenDJ, as described in the *OpenDJ Installation Guide*.

During installation, populate OpenDJ with the data in the `Example.ldif` file, available in the `sample6/data` directory.

The actions run in this sample should not change any data on the AD server.

3.7.2. Sample 6 with a Simulated AD Deployment

In this sample, an AD deployment is simulated with a different baseDN (`dc=fakead,dc=com`) on the same OpenDJ server instance. You can also simulate the AD server with a separate OpenDJ instance, running on the same host, as long as the two instances communicate on different ports. The data

for the simulated AD instance is contained in the file `AD.ldif`. The data for the OpenDJ instance is contained in the file `Example.ldif`.

3.7.3. External Configuration

3.7.3.1. Prepare OpenDJ For LiveSync

This sample assumes a replicated OpenDJ server on the localhost system. When configured, OpenDJ replication includes an External Change Log (ECL), required to support liveSync. LiveSync detects changes in OpenDJ by reading the ECL.

Follow these steps to install and configure an OpenDJ instance.

1. Download OpenDJ from ForgeRock's BackStage site and extract the zip archive.
2. Set up OpenDJ and import the data file for this sample, as follows:

```
$ cd /path/to/opendj
$ ./setup \
  directory-server
  \
  --hostname localhost
  \
  --ldapPort 1389
  \
  --rootUserDN "cn=Directory Manager"
  \
  --rootUserPassword password
  \
  --adminConnectorPort 4444
  \
  --baseDN dc=com
  \
  --ldifFile /path/to/openidm/samples/sample6/data/Example.ldif
  \
  --acceptLicense
Validating parameters .....Done.
Configuring Certificates .....Done.
Configuring server .....Done.
Importing data from /path/to/openidm/samples/sample6/data/Example.ldif .....Done.
Starting Directory Server .....Done.

To see basic server status and configuration, you can launch /path/to/opendj/bin/status
```

The sample assumes the following configuration:

- The OpenDJ server is installed on the localhost.
- The server listens for LDAP connections on port 1389.
- The administration connector port is 4444.
- The root user DN is `cn=Directory Manager`.

- The root user password is `password`.
3. Configure the OpenDJ directory server as a replication server.

To enable liveSync, this server must be configured for replication, even if it does not actually participate in a replication topology. The following commands configure the server for replication.

```
$ cd /path/to/opendj/bin
./dsconfig create-replication-server \
--hostname localhost \
--port 4444 \
--bindDN "cn=Directory Manager" \
--bindPassword password \
--provider-name "Multimaster Synchronization" \
--set replication-port:8989 \
--set replication-server-id:2 \
--trustAll \
--no-prompt

$ ./dsconfig create-replication-domain \
--hostname localhost \
--port 4444 \
--bindDN "cn=Directory Manager" \
--bindPassword password \
--provider-name "Multimaster Synchronization" \
--domain-name fakead_com \
--set base-dn:dc=fakead,dc=com \
--set replication-server:localhost:8989 \
--set server-id:3 \
--trustAll \
--no-prompt
```

Once OpenDJ is configured, you can proceed with either an active or simulated AD deployment.

3.7.3.2. External Configuration for an Active AD Deployment

To configure an active AD deployment for sample 6, open the `provisioner.openicf-realad.json` file in a text editor. Update it as needed. At minimum, you should check and if needed update the following parameters in that file, as shown in the following table:

Table 3.1. Key Parameters for an Active AD Configuration

Option	Description
host	The hostname/IP address of the AD server
port	The LDAP port; the default is 389.
ssl	By default, SSL is not used.
principal	The full DN of the account to bind with, such as "CN=Administrator,CN=Users,DC=example,DC=com"

Option	Description
credentials	If a password is used, replace null with that password. When OpenIDM starts, it encrypts that password in the <code>provisioner.openicf-realad.conf</code> file.
baseContexts	The DN's for account containers, such as ["CN=Users,DC=Example,DC=com"]
baseContextsToSynchronize	Set to the same value as <code>baseContexts</code>
accountSearchFilter	Default searches for active user (not computer) accounts
accountSynchronizationFilter	Default synchronizes with active user (not computer) accounts

If you do not want to filter out computer and disabled user accounts, set the `accountSearchFilter` and `accountSynchronizationFilter` to `null`.

3.7.3.3. External Configuration for a Simulated AD Deployment

Not everyone has a testable instance of AD readily available. For such administrators, you can use the `AD.ldif` file from the `data/` subdirectory to simulate an AD deployment in a separate suffix on your OpenDJ directory server.

If you have not already done so, copy the `provisioner.openicf-fakead.json` file to the `conf` subdirectory.

As previously mentioned, you can use a separate OpenDJ instance to simulate the AD server. However, the following instructions assume that the simulated AD server runs on the same OpenDJ instance.

Open the `provisioner.openicf-fakead.json` file and note the following:

- OpenDJ directory server uses port 1389 by default for users who cannot use privileged ports, so this is the port that is specified in the provisioner file. Adjust the port if your OpenDJ server is listening on a different port.
- The simulated AD server uses the base DN `dc=fakead,dc=com`.

To load the data for the simulated AD instance, launch the OpenDJ control panel (`/path/to/opensharedir/bin/control-panel`), add the base DN (`dc=fakead,dc=com`), then import the `sample6/data/AD.ldif` file. When you import the `AD.ldif` file, select "Append to Existing Data", not "Overwrite Existing Data". Otherwise, the data in the `dc=example,dc=com` base DN will be overwritten.

Alternatively, you could run the following command:

```
$ cd /path/to/openssl/bin
$ ./ldapmodify
\
--bindDN "cn=Directory Manager"
\
--bindPassword password
\
--hostname localhost
\
--port 1389
\
--filename /path/to/openidm/samples/sample6/data/AD.ldif
```

3.7.4. Start OpenIDM with Sample 6

Now that OpenDJ and a real or simulated AD database is configured, prepare OpenIDM as described in Section 1.3, "Preparing the Server". You can then start OpenIDM with the configuration for sample 6.

```
$ cd /path/to/openidm
$ ./startup.sh -p samples/sample6
```

3.7.5. Run the Sample

The following sections show how to run the sample with command-based reconciliation with a REST call, and to configure scheduled reconciliation with liveSync.

3.7.5.1. Using Reconciliation

Now that OpenIDM is in operation, review the entries in the OpenDJ data store. When you run reconciliation, any entries that share the same `uid` with the AD data store will be updated with the contents from AD.

If you have set up the simulated AD data store as described in Section 3.7.3.3, "External Configuration for a Simulated AD Deployment", compare the entries for `uid=jdoe` as shown in the `AD.ldif` and `Example.ldif` files. Note the different values of `givenName` for `uid=jdoe`.

Run reconciliation over the REST interface. If you have followed the instructions for the simulated AD data store, the following command takes the information for user `jdoe` imported from the `AD.ldif` file, with a `givenName` of Johnny, and synchronizes that information to the LDAP database, overwriting the `givenName` of John for that same user `jdoe`.

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request POST \
"http://localhost:8080/openidm/recon?
_action=recon&mapping=systemAdAccounts_managedUser&waitForCompletion=true"
```

The reconciliation operation returns a reconciliation run ID, and the status of the operation.

```
{
  "state": "SUCCESS",
  "_id": "985ee939-fbe1-4607-a757-00b404b4ef77"
}
```

The reconciliation operation synchronizes the data in the AD server with the OpenIDM repository (managed/user). That information is then automatically synchronized to the OpenDJ server, as described in Section 14.13, "Synchronization Situations and Actions" in the *Integrator's Guide*.

After reconciliation, list all users in the OpenDJ server data store.

```
$ curl \
  --header "X-OpenIDM-Username: openidm-admin" \
  --header "X-OpenIDM-Password: openidm-admin" \
  --request GET \
  "http://localhost:8080/openidm/system/ldap/account?_queryId=query-all-ids"
```

The result should resemble the following JSON object.

```
{
  "result": [ {
    "dn" : "uid=jdoe,ou=People,dc=example,dc=com",
    "_id" : "uid=jdoe,ou=People,dc=example,dc=com"
  }, {
    "dn" : "uid=bjensen,ou=People,dc=example,dc=com",
    "_id" : "uid=bjensen,ou=People,dc=example,dc=com"
  } ],
  ...
}
```

You see only two entries, as the `uid=jdoe` entry from `dc=fakead,dc=com` overwrites the original LDAP entry for `uid=jdoe` in the reconciled LDAP data store.

To read the user object in the OpenDJ server, run the `ldapsearch` command. The following example returns the entry for user `uid=jdoe`:

```
$ ./ldapsearch \
  --port 1389 \
  --baseDN dc=example,dc=com \
  "(uid=jdoe)"
```

3.7.5.2. Using LiveSync

You can trigger a reconciliation operation by configuring a schedule, or by launching the operation directly over the REST interface. You can also launch a liveSync operation over REST, but liveSync requires a configured schedule to poll for changes. When this sample's default liveSync schedule (`schedule-activeSynchroniser_systemAdAccount.json`) is enabled, a liveSync operation is launched every 15 seconds.

LiveSync pushes changes made in the AD data store to the OpenIDM repository, automatically.

The liveSync schedule is disabled by default. To activate liveSync, change the value of the `enabled` property from `false` to `true`.

```
{
  "enabled" : false,
  "type" : "cron",
  "schedule" : "0/15 * * * * ?",
  "persisted" : true,
  "invokeService" : "provisioner",
  "invokeContext" : {
    "action" : "liveSync",
    "source" : "system/ad/account"
  },
  "invokeLogLevel" : "debug"
}
```

Procedure 3.9. Testing LiveSync

Now you can test liveSync. This procedure assumes that you have configured OpenDJ using the parameters and commands described in this section.

1. Create an LDIF file with a new user entry (`uid=bsmith`) that will be added to the simulated AD data store.
2. The following is the contents of a sample `bsmith.ldif` file for demonstration purposes:

```
dn: uid=bsmith,ou=People,dc=fakead,dc=com
objectClass: person
objectClass: inetOrgPerson
objectClass: organizationalPerson
objectClass: top
givenName: Barry
description: Created to see liveSync work
uid: bsmith
cn: Barry
sn: Smith
mail: bsmith@example.com
telephoneNumber: 1-415-523-0772
userPassword: passw0rd
```

3. Navigate to the `/path/to/openssl/bin` directory.
4. Use the `ldapmodify` command to add the `bsmith.ldif` file to the directory.

```
$ ./ldapmodify \
--port 1389 \
--bindDN "cn=Directory Manager" \
--bindPassword password \
--filename /path/to/bsmith.ldif
```

5. Now you can test synchronization by viewing the new user in the OpenIDM repository. The easiest way to do this, is through OpenIDM UI. You should be able to log into the UI with any of the accounts in the AD data store. For this example, log into the UI as user `bsmith`, with password `passw0rd`. The fact that you can log into the UI as this new user indicates that liveSync has synchronized the user from the AD data store to the managed/user repository.

6. Implicit synchronization pushes this change out to the OpenDJ data store. To test this synchronization operation, search the OpenDJ baseDN for the new user entry.

```
$ ./ldapsearch \  
--port 1389 \  
--baseDN ou=people,dc=example,dc=com \  
"(uid=bsmith)"
```

3.8. Linking Historical Accounts

This sample demonstrates the retention of inactive (historical) LDAP accounts that have been linked to a corresponding managed user account. The sample builds on sample 2b and uses the LDAP connector to connect to an OpenDJ instance. You can use any LDAP-v3 compliant directory server.

In this sample, OpenIDM is the source resource. Managed users in the OpenIDM repository maintain a list of the accounts that they have been linked to on the local LDAP server. This list is stored in the `historicalAccounts` field of the managed user entry. The list contains a reference to all past and current LDAP accounts. Each LDAP account in the list is represented as a *relationship* and includes information about the date the accounts were linked or unlinked, and whether the account is currently active. For more information about relationship objects, see Section 9.5, "Managing Relationships Between Objects" in the *Integrator's Guide*.

This sample includes the following custom scripts, in its `script` directory:

- `onLink-managedUser_systemLdapAccounts.js`

When a managed user object is linked to a target LDAP object, this script creates the relationship entry in the managed user's `historicalAccounts` property. The script adds two relationship properties:

- `linkDate` — specifies the date that the link was created.
 - `active` — boolean true/false. When set to true, this property indicates that the target object is *currently* linked to the managed user account.
- `onUnlink-managedUser_systemLdapAccounts.js`

When a managed user object is unlinked from a target LDAP object, this script updates that relationship entry's properties with an `unlinkDate` that specifies when the target was unlinked, and sets the `active` property to false, indicating that the target object is no longer linked.

- `check_account_state_change.js`

During liveSync or reconciliation, this script checks if the LDAP account state has changed. If the state has changed, the script updates the historical account properties to indicate the new state (enabled or disabled), and the date that the state was changed. This date can only be approximated and is set to the time that the change was detected by the script.

- `ldapBackCorrelationQuery.js`

This script correlates entries in the LDAP directory with managed user identities in OpenIDM.

3.8.1. Configuring the LDAP Server

This sample expects the configuration for the external LDAP server to be the same as described in Section 3.1.1, "LDAP Server Configuration".

The following steps provide setup instructions for OpenDJ director server version 5. Adjust these instructions if you are using an older version of OpenDJ, or an alternative LDAP server.

The LDIF data for this sample is provided in the file `openidm/samples/historicalaccountlinking/data/Example.ldif`. Import this data during your OpenDJ setup.

Although there is only one LDAP server in this example, you must enable *replication* on that server, so that the server has an external change log. The change log is required for liveSync between OpenDJ and OpenIDM.

1. Download OpenDJ from ForgeRock's BackStage site and extract the zip archive.
2. Set up OpenDJ and import the `Example.ldif` file for this sample:

```
$ cd /path/to/opendj
$ ./setup \
  directoryserver \
  --hostname localhost \
  --ldapPort 1389 \
  --rootUserDN "cn=Directory Manager" \
  --rootUserPassword password \
  --adminConnectorPort 4444 \
  --baseDN dc=com \
  --ldifFile /path/to/openidm/samples/historicalaccountlinking/data/Example.ldif \
  --acceptLicense
.Validating parameters .....Done.
Configuring Certificates .....Done.
Configuring server .....Done.
Importing data from /path/to/openidm/samples/historicalaccountlinking/data/Example.ldif .....Done.
Starting Directory Server .....Done.
```

To see basic server status and configuration, you can launch `/path/to/opendj/bin/status`

3. Enable replication:

```
$ cd /path/to/openssh/bin
./dsconfig create-replication-server \
--hostname localhost \
--port 4444 \
--bindDN "cn=Directory Manager" \
--bindPassword password \
--provider-name "Multimaster Synchronization" \
--set replication-port:8989 \
--set replication-server-id:2 \
--trustAll \
--no-prompt
$ bin/dsconfig create-replication-domain \
--hostname localhost \
--port 4444 \
--bindDN "cn=Directory Manager" \
--bindPassword password \
--provider-name "Multimaster Synchronization" \
--set base-dn:dc=example,dc=com \
--set replication-server:localhost:8989 \
--set server-id:3 \
--domain-name example_com \
--trustAll \
--no-prompt
```

3.8.2. Running the Historical Accounts Sample

This section walks you through each step of the sample to demonstrate how historical accounts are stored.

1. Prepare OpenIDM as described in Section 1.3, "Preparing the Server", then start OpenIDM with the configuration for the historical accounts sample:

```
$ cd /path/to/openidm
$ ./startup.sh -p samples/historicalaccountlinking/
Executing ./startup.sh...
Using OPENIDM_HOME: /path/to/openidm
Using PROJECT_HOME: /path/to/openidm/samples/historicalaccountlinking/
Using OPENIDM_OPTS: -Xmx1024m -Xms1024m
Using LOGGING_CONFIG: -Djava.util.logging.config.file=
/path/to/openidm/samples/historicalaccountlinking//conf/logging.properties
Using boot properties at
/path/to/openidm/samples/historicalaccountlinking/conf/boot/boot
.properties
-> OpenIDM ready
```

2. Create a user, Joe Smith, in OpenIDM.

The following command creates the user over REST, and assigns the user the ID `joesmith`:


```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--header "Content-Type: application/json" \
--request POST \
--data '{
  "userName": "joe.smith",
  "givenName": "Joe",
  "sn": "Smith",
  "password": "Passw0rd",
  "displayName": "Joe Smith",
  "mail": "joe.smith@example.com",
  "_id": "joesmith"
}' \
"http://localhost:8080/openidm/managed/user?_action=create"
{
  "_id": "joesmith",
  "_rev": "1",
  "userName": "joe.smith",
  "givenName": "Joe",
  "sn": "Smith",
  "displayName": "Joe Smith",
  "mail": "joe.smith@example.com",
  "accountStatus": "active",
  "effectiveRoles": [],
  "effectiveAssignments": []
}
```

3. Verify that the user Joe Smith was created in OpenDJ.

Because implicit synchronization is enabled by default, any change to the managed/user repository should be propagated to OpenDJ. For more information about implicit synchronization, see Section 14.1, "Types of Synchronization" in the *Integrator's Guide*.

The following command returns all IDs in OpenDJ and shows that user joesmith was created successfully:

```

$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/system/ldap/account?_queryId=query-all-ids"
{
  "result": [
    {
      "_id": "uid=jdoe,ou=People,dc=example,dc=com",
      "dn": "uid=jdoe,ou=People,dc=example,dc=com"
    },
    {
      "_id": "uid=bjensen,ou=People,dc=example,dc=com",
      "dn": "uid=bjensen,ou=People,dc=example,dc=com"
    },
    {
      "_id": "uid=joe.smith0,ou=People,dc=example,dc=com",
      "dn": "uid=joe.smith0,ou=People,dc=example,dc=com"
    }
  ]
  ,
  ...
}

```

Note that Joe Smith's `uid` in OpenDJ is appended with a `0`. The `onCreate` script, defined in the mapping (`sync.json`), increments the `uid` each time a new OpenDJ entry is linked to the same managed user object.

4. Verify that the historical account *relationship object* that corresponds to this linked LDAP account was created in the OpenIDM repository.

The following command returns all of the `historicalAccounts` for user joesmith:

```

$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/managed/user/joesmith/historicalAccounts?_queryId=query-all"
{
  "result": [
    {
      "_ref": "system/ldap/account/uid=joe.smith0,ou=People,dc=example,dc=com",
      "_refProperties": {
        "stateLastChanged": "Mon Nov 30 2015 14:45:22 GMT-0800 (PST)",
        "state": "enabled",
        "active": true,
        "linkDate": "Mon Nov 30 2015 14:45:22 GMT-0800 (PST)",
        "_id": "ff6913ce-a252-4dc9-a060-b8b56bb32bf4",
        "_rev": "1"
      }
    }
  ]
  ,
  ...
}

```

At this stage, Joe Smith has only one historical account link — the link to `uid=joe.smith0,ou=People,dc=example,dc=com`. Note that the relationship properties (`_refProperties`) show the following information about the linked accounts:

- The date on which the accounts were linked
 - The fact that this link is currently active
 - The state of the account in OpenDJ (`enabled`)
5. Enable the liveSync schedule to propagate changes made in OpenDJ to the managed user repository.

To start liveSync, set `enabled` to `true` in the `conf/schedule-liveSync.json` file:

```
$ cd /path/to/openidm
$ more samples/historicalaccountlinking/conf/schedule-liveSync.json
{
  "enabled" : true,
  "type" : "cron",
  "schedule" : "0/15 * * * * ?"
  ,
  ...
}
```

6. Use the `manage-account` command in the `opendj/bin` directory to disable Joe Smith's account in OpenDJ:

```
$ cd /path/to/opendj
$ bin/manage-account set-account-is-disabled
\
--port 4444
\
--bindDN "cn=Directory Manager"
\
--bindPassword password
\
--operationValue true
\
--targetDN uid=joe.smith0,ou=people,dc=example,dc=com
\
--trustAll
Account Is Disabled: true
```

Within 15 seconds, according to the configured schedule, liveSync should pick up the change. OpenIDM should then adjust the `state` property in Joe Smith's managed user account.

7. Check Joe Smith's historical accounts again, to make sure that the state of this linked account has changed:

```

$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/managed/user/joesmith/historicalAccounts?_queryId=query-all"
{
  "result": [
    {
      "_ref": "system/ldap/account/uid=joe.smith0,ou=People,dc=example,dc=com",
      "_refProperties": {
        "stateLastChanged": "Mon Nov 30 2015 14:54:45 GMT-0800 (PST)",
        "state": "disabled",
        "active": true,
        "linkDate": "Mon Nov 30 2015 14:45:22 GMT-0800 (PST)",
        "_id": "fff6913ce-a252-4dc9-a060-b8b56bb32bf4",
        "_rev": "2"
      }
    }
  ]
}
...
}
    
```

- Now, deactivate Joe Smith's managed user account by setting his `accountStatus` property to inactive.

You can deactivate the account over the REST interface, or by using the Admin UI.

To use the Admin UI, simply select Manage > User, select Joe Smith's account and change his Status to inactive, on his Details tab.

The following command deactivates Joe Smith's account over REST:

```

$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--header "Content-Type: application/json" \
--request PATCH \
--data '[
  { "operation" : "replace",
    "field" : "accountStatus",
    "value" : "inactive" }
]' \
"http://localhost:8080/openidm/managed/user/joesmith"
{
  "_id": "joesmith",
  "_rev": "3",
  "userName": "joe.smith",
  ...
  "accountStatus": "inactive",
  ...
}
    
```

- Request Joe Smith's historical accounts again:

```

$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/managed/user/joesmith/historicalAccounts?_queryId=query-all"
{
  "result": [
    {
      "_ref": "system/ldap/account/uid=joe.smith0,ou=People,dc=example,dc=com",
      "_refProperties": {
        "stateLastChanged": "Mon Nov 30 2015 14:54:45 GMT-0800 (PST)",
        "state": "disabled",
        "active": false,
        "linkDate": "Mon Nov 30 2015 14:45:22 GMT-0800 (PST)",
        "unlinkDate": "Mon Nov 30 2015 14:58:30 GMT-0800 (PST)",
        "_id": "fff6913ce-a252-4dc9-a060-b8b56bb32bf4",
        "_rev": "3"
      }
    }
  ]
  ,
  ...
}

```

10. Activate Joe Smith's managed user account by setting his `accountStatus` property to active. This action should create a new entry in OpenDJ (with `uid=joe.smith1`), and a new link from Joe Smith's managed user object to that OpenDJ entry.

You can activate the account over the REST interface, or by using the Admin UI, as described previously.

The following command activates Joe Smith's account over REST:

```

$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--header "Content-Type: application/json" \
--request PATCH \
--data '[
  { "operation" : "replace",
    "field" : "accountStatus",
    "value" : "active" }
]' \
"http://localhost:8080/openidm/managed/user/joesmith"
{
  "_id": "joesmith",
  "_rev": "4",
  "userName": "joe.smith",
  ...
  "accountStatus": "active",
  ...
}

```

11. Verify that a new LDAP entry for user Joe Smith was created in OpenDJ.

The following command returns all IDs in OpenDJ and shows that two OpenDJ entries for Joe Smith `uid=joe.smith0` and `uid=joe.smith1`.

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/system/ldap/account?_queryId=query-all-ids"
{
  "result": [
    {
      "_id": "uid=jdoe,ou=People,dc=example,dc=com",
      "dn": "uid=jdoe,ou=People,dc=example,dc=com"
    },
    {
      "_id": "uid=bjensen,ou=People,dc=example,dc=com",
      "dn": "uid=bjensen,ou=People,dc=example,dc=com"
    },
    {
      "_id": "uid=joe.smith0,ou=People,dc=example,dc=com",
      "dn": "uid=joe.smith0,ou=People,dc=example,dc=com"
    },
    {
      "_id": "uid=joe.smith1,ou=People,dc=example,dc=com",
      "dn": "uid=joe.smith1,ou=People,dc=example,dc=com"
    }
  ]
  ,
  ...
}
```

12. Request Joe Smith's historical accounts again:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/managed/user/joesmith/historicalAccounts?_queryId=query-all"
{
  "result": [
    {
      "_ref": "system/ldap/account/uid=joe.smith0,ou=People,dc=example,dc=com",
      "_refProperties": {
        "stateLastChanged": "Mon Nov 30 2015 14:54:45 GMT-0800 (PST)",
        "state": "disabled",
        "active": false,
        "linkDate": "Mon Nov 30 2015 14:45:22 GMT-0800 (PST)",
        "unlinkDate": "Mon Nov 30 2015 14:58:30 GMT-0800 (PST)",
        "_id": "fff6913ce-a252-4dc9-a060-b8b56bb32bf4",
        "_rev": "3"
      }
    },
    {
      "_ref": "system/ldap/account/uid=joe.smith1,ou=People,dc=example,dc=com",
      "_refProperties": {
        "stateLastChanged": "Mon Nov 30 2015 15:00:00 GMT-0800 (PST)",
        "state": "enabled",
      }
    }
  ]
}
```

```
"active": true,  
"linkDate": "Mon Nov 30 2015 15:00:00 GMT-0800 (PST)",  
"_id": "08443775-7420-4994-bf86-9b29a986bfc9",  
"_rev": "1"  
  }  
},  
  ],  
  ...  
}
```

Note that Joe Smith's entry now shows two OpenDJ accounts, but that only the link to `uid=joe.smith1` is `enabled` and `active`.

3.9. Storing Multiple Passwords For Managed Users

This sample demonstrates how to set up multiple passwords for managed users and how to synchronize separate passwords to different external resources. The sample assumes the following scenario:

- The managed/user repository is the source system.
- There are two target LDAP servers — `ldap` and `ldap2`.

For the purposes of this sample, the two servers are represented by two separate organizational units on a single OpenDJ instance.

- Managed user objects have two additional password fields, each mapped to one of the two LDAP servers.
- The two LDAP servers have different requirements for password policy and encryption.
- Both LDAP servers have a requirement for a password history policy, but the history size differs for the two policies.

The sample shows how to extend the password history policy, described in [Section 17.1.1, "Creating a Password History Policy"](#) in the *Integrator's Guide*, to apply to multiple password fields.

- The value of a managed user's `password` field is used by default for the additional passwords *unless* the CREATE, UPDATE, or PATCH requests on the managed user explicitly contain a value for these additional passwords.

The sample includes the following customized configuration files in the `conf` directory:

- `provisioner.openicf-ldap.json` configures the first LDAP connection.
- `conf/provisioner.openicf-ldap2.json` configures the second LDAP connection.
- `sync.json` provides the mappings from the OpenIDM managed user repository to the respective LDAP servers.

- `managed.json` contains a customized schema for managed users that includes the additional password fields.

For details of the customizations to these configuration files, see the [README](#) provided with the sample.

The sample includes the following customized scripts in the `script` directory:

- `onCreate-onUpdate-sync.js` performs custom mapping logic. Specifically, this script maps the pre-hashed password value and sets the target object DN on create events.
- `storeFields.groovy` stores the pre-hashed values of fields in the context chain, on validate events.
- `onCreate-user-custom.js` and `onUpdate-user-custom.js` are used for validation of the password history policy when a user is created or updated.
- `pwpolicy.js` is an additional policy script for the password history policy.
- `set-additional-passwords.js` populates the values of the additional password fields with the value of the main `password` field if the additional fields are not included in the request content.

Note

This sample does not support creation of new users in the Admin UI.

3.9.1. Understanding the Password History Policy

The sample includes a custom password history policy. Although the sample is concerned only about the history of passwords, you can use this policy to enforce history validation on any managed object property.

The following configuration changes set up the password history policy:

- A `fieldHistory` property is added to managed users. The value of this field is a map of field names to a list of historical values for that field. These lists of values are used by the policy to determine if a new value has previously been used.
- The `onCreate-user-custom.js` script performs the standard `onCreate` tasks for a managed user object but also stores the initial value of each of the fields that OpenIDM must keep a history of. The script is passed the following configurable properties:
 - `historyFields` — a list of the fields to store history on.
 - `historySize` — the number of historical fields to store.
- The `onUpdate-user-custom.js` compares the old and new values of the historical fields on update events, to determine if the values have changed. When a new value is detected, it is stored in the list of historical values for that field.

This script also contains logic to deal with the comparison of encrypted or hashed field values. The script is passed the following configurable properties:

- `historyFields` — a list of the fields to store history on.
- `historySize` — the number of historical fields to store.
- The `pwpolicy.js` script contains the additional policy definition for the historical password policy. This script compares the new field value with the values in the list of historical values for each field.

The `policy.json` configuration includes this script in its `additionalFiles` list, so that the policy service loads the new policy definition. The new policy takes a `historyLength` parameter, which indicates the number of historical values to enforce the policy on. This number must not exceed the `historySize` specified in the `onCreate` and `onUpdate` scripts.

- The `ldapPassword` and `ldap2Password` fields in the managed user schema have been updated with the policy. For the purposes of this sample the `historySize` has been set to 2 for `ldapPassword` and to 4 for `ldap2Password`.

3.9.2. Configuring the LDAP Server

This sample expects the configuration for the external LDAP server to be the same as described in Section 3.1.1, "LDAP Server Configuration".

The following steps provide setup instructions for OpenDJ director server version 5. Adjust these instructions if you are using an older version of OpenDJ, or an alternative LDAP server.

The LDIF data for this sample is provided in the file `openidm/samples/multiplepasswords/data/Example.ldif`. Import this data during your OpenDJ setup.

1. Download OpenDJ from ForgeRock's BackStage site and extract the zip archive.
2. Set up OpenDJ and import the `Example.ldif` file for this sample:

```
$ cd /path/to/opendj
$ ./setup \
  directoryserver \
  --hostname localhost \
  --ldapPort 1389 \
  --rootUserDN "cn=Directory Manager" \
  --rootUserPassword password \
  --adminConnectorPort 4444 \
  --baseDN dc=com \
  --ldifFile /path/to/openidm/samples/multiplepasswords/data/Example.ldif \
  --acceptLicense
.Validating parameters .....Done.
Configuring Certificates .....Done.
Configuring server .....Done.
Importing data from /path/to/openidm/samples/multiplepasswords/data/Example.ldif .....Done.
Starting Directory Server .....Done.
```

To see basic server status and configuration, you can launch `/path/to/opendj/bin/status`

3. Run an `ldapsearch` on the LDAP directory and look at the organizational units:

```
$ cd /path/to/openssl
$ bin/ldapsearch
\
--hostname localhost
\
--port 1389
\
--bindDN "cn=directory manager"
\
--bindPassword password
\
--baseDN "dc=example,dc=com" \
"ou=*" \
ou
dn: ou=People,dc=example,dc=com
ou: people

dn: ou=Customers,dc=example,dc=com
ou: people
ou: Customers
```

The organizational units, `ou=People` and `ou=Customers`, represent the two different target LDAP systems that our mappings point to.

3.9.3. Demonstrating the Use of Multiple Accounts

This section starts OpenIDM with the sample configuration, then creates a user with multiple passwords, adhering to the different policies in the configured password policy. The section tests that the user was synchronized to two separate LDAP directories, with the different required passwords, and that the user can bind to each of these LDAP directories.

1. Prepare OpenIDM as described in Section 1.3, "Preparing the Server", then start OpenIDM with the configuration for the multiple passwords sample.

```
$ cd /path/to/openidm
$ ./startup.sh -p samples/multiplepasswords
```

2. Create a user in OpenIDM. Include a main `password` field but no additional password fields. The additional password fields (`ldapPassword` and `ldap2Password`) will be populated with the value of the main `password` field as a result of the script described previously.

For the purposes of this example, we will not use the Admin UI, so that the result of each command can be clearly seen. Create the user over REST, by running the following command:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--header "Content-Type: application/json" \
--request PUT \
--data '{
  "userName": "jdoe",
  "givenName": "John",
```

```
"sn" : "Doe",
"displayname" : "John Doe",
"mail" : "john.doe@example.com",
"password" : "Passw0rd"
}, \
"http://localhost:8080/openidm/managed/user/jdoe"
{
  "code": 403,
  "reason": "Forbidden",
  "message": "Policy validation failed",
  "detail": {
    "result": false,
    "failedPolicyRequirements": [
      {
        "policyRequirements": [
          {
            "policyRequirement": "AT_LEAST_X_CAPITAL_LETTERS",
            "params": {
              "numCaps": 2
            }
          }
        ]
      },
      {
        "policyRequirements": [
          {
            "policyRequirement": "AT_LEAST_X_NUMBERS",
            "params": {
              "numNums": 2
            }
          }
        ]
      }
    ],
    "property": "ldap2Password"
  }
}
]
```

Notice that the create request failed with a policy validation failure on the two new password fields. Although the password met the requirement for the main `password` field, the user could not be created because the password did not meet the requirements of the `ldapPassword` and `ldap2Password` fields.

You can fix this problem either by updating the `password` field to one that passes both of the new requirements, or by updating the individual password fields to specifically pass their individual validation requirements.

3. Now, try to create user `jdoe` again, this time providing individual values for each of the different password fields, that comply with the three different password policies:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--header "Content-Type: application/json" \
--request PUT \
--data '{
```

```

"userName": "jdoe",
"givenName": "John",
"sn": "Doe",
"displayName": "John Doe",
"mail": "john.doe@example.com",
"password": "Passw0rd",
"ldapPassword": "PPassw0rd",
"ldap2Password": "Passw00rd"
}, \
"http://localhost:8080/openidm/managed/user/jdoe"
{
  "_id": "jdoe",
  "_rev": "1",
  "userName": "jdoe",
  "givenName": "John",
  "sn": "Doe",
  "displayName": "John Doe",
  "mail": "john.doe@example.com",
  "ldapPassword": {
    "$crypto": {
      "value": {
        "algorithm": "SHA-256",
        "data": "CpbVZlXAEFL/LUqWyq9Bcks/tLVwJ0pHrc/smLwf8UmC/0BDtEKRo1o2IjB6mNFz"
      },
      "type": "salted-hash"
    }
  },
  "ldap2Password": {
    "$crypto": {
      "value": {
        "iv": "TbJlRF+cSFe0guclh8AZVg==",
        "data": "zQ250CXfr3QJ3cBKjpCQhQ==",
        "cipher": "AES/CBC/PKCS5Padding",
        "key": "openidm-sym-default"
      },
      "type": "x-simple-encryption"
    }
  }
}
,
...
}

```

The user has been created with three different passwords that comply with three distinct password policies. The passwords have been hashed or encrypted, as defined in the `managed.json` file.

4. As a result of implicit synchronization, two separate LDAP accounts should have been created for user `jdoe` on our two simulated LDAP servers. For more information about implicit synchronization, see Section 14.1, "Types of Synchronization" in the *Integrator's Guide*.
5. Query the IDs in the LDAP directory as follows:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/system/ldap/account?_queryId=query-all-ids"
{
  "result" : [ {
    "_id" : "uid=jdoe,ou=People,dc=example,dc=com",
    "dn" : "uid=jdoe,ou=People,dc=example,dc=com"
  }, {
    "_id" : "uid=jdoe,ou=Customers,dc=example,dc=com",
    "dn" : "uid=jdoe,ou=Customers,dc=example,dc=com"
  } ]
,
...
}
```

Note that jdoe has two entries - one in **ou=People** and one in **ou=Customers**.

- Now, see if you can search each LDAP server, as user jdoe, with the separate passwords that you created for each directory.

This step will indicate that the passwords were propagated correctly to the separate LDAP servers.

```
$ cd /path/to/opendj
$ bin/ldapsearch
\
--hostname localhost
\
--port 1389
\
--bindDN uid=jdoe,ou=People,dc=example,dc=com
\
--bindPassword PPassw0rd
\
--baseDN dc=example,dc=com \
uid=jdoe
dn: uid=jdoe,ou=People,dc=example,dc=com
objectClass: organizationalPerson
objectClass: person
objectClass: inetOrgPerson
objectClass: top
uid: jdoe
mail: john.doe@example.com
sn: Doe
cn: John Doe
userPassword: {SSHA}ot11NT7zidSxXEDtNE+8qQjyfIE3CDbywKGYmQ==
givenName: John

dn: uid=jdoe,ou=Customers,dc=example,dc=com
objectClass: organizationalPerson
objectClass: person
objectClass: inetOrgPerson
objectClass: top
uid: jdoe
mail: john.doe@example.com
```

```

sn: Doe
cn: John Doe
givenName: John
$ bin/ldapsearch
\
--hostname localhost
\
--port 1389
\
--bindDN uid=jdoe,ou=Customers,dc=example,dc=com
\
--bindPassword Passw00rd
\
--baseDN dc=example,dc=com \
uid=jdoe
dn: uid=jdoe,ou=People,dc=example,dc=com
objectClass: organizationalPerson
objectClass: person
objectClass: inetOrgPerson
objectClass: top
uid: jdoe
mail: john.doe@example.com
sn: Doe
cn: John Doe
userPassword: {SSHA}ot11NT7zidSxXEDtNE+8qQjyfIE3CDbywKGYmQ==
givenName: John

dn: uid=jdoe,ou=Customers,dc=example,dc=com
objectClass: organizationalPerson
objectClass: person
objectClass: inetOrgPerson
objectClass: top
uid: jdoe
mail: john.doe@example.com
sn: Doe
cn: John Doe
givenName: John
  
```

7. Patch jdoe's managed user entry to change his `ldapPassword`.

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--header "Content-Type: application/json" \
--request PATCH \
--data '[ {
  "operation" : "replace",
  "field" : "ldapPassword",
  "value" : "TTestw0rd"
} ]' \
"http://localhost:8080/openidm/managed/user/jdoe"
{
  "_id": "jdoe",
  "_rev": "2",
  "userName": "jdoe",
  "givenName": "John",
  "sn": "Doe",
  "displayName": "John Doe",
  ...
  "ldapPassword": {
    "$crypto": {
      "value": {
        "algorithm": "SHA-256",
        "data": "Vc6hvmzXaSSdG9Wroq0Tg3PQVdixhpg9tD/uKT610Z/H5iC6vso0pE0/R5FaiDUg"
      },
      "type": "salted-hash"
    }
  },
  ...
}
```

8. Search the "first" LDAP server again, as user jdoe, with this new password to verify that the password change was propagated correctly to the LDAP server.

```

$ cd /path/to/opendj
$ bin/ldapsearch
\
--hostname localhost
\
--port 1389
\
--bindDN uid=jdoe,ou=People,dc=example,dc=com
\
--bindPassword TTestw0rd
\
--baseDN dc=example,dc=com \
uid=jdoe
dn: uid=jdoe,ou=People,dc=example,dc=com
objectClass: organizationalPerson
objectClass: person
objectClass: inetOrgPerson
objectClass: top
userPassword: {SSHA}pXV9/eZq6L30L/1GTsMV/39Dzjv/zHqIhWpLRw==
uid: jdoe
mail: john.doe@example.com
sn: Doe
givenName: John
cn: John Doe

dn: uid=jdoe,ou=Customers,dc=example,dc=com
objectClass: organizationalPerson
objectClass: person
objectClass: inetOrgPerson
objectClass: top
uid: jdoe
mail: john.doe@example.com
sn: Doe
cn: John Doe
givenName: John

```

3.9.4. Demonstrating the Use of the Password History Policy

This section patches managed user jdoe's entry, changing his `ldapPassword` a number of times, to demonstrate the application of the password history policy.

1. Send the following patch requests, changing the value of jdoe's `ldapPassword` each time:

```

$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--header "Content-Type: application/json" \
--request PATCH \
--data '[ {
  "operation" : "replace",
  "field" : "ldapPassword",
  "value" : "TTestw0rd1"
} ]' \
"http://localhost:8080/openidm/managed/user/jdoe"
{
  "_id": "jdoe",

```



```

    "_rev": "3",
    "userName": "jdoe",
    "givenName": "John",
    "sn": "Doe",
    "displayName": "John Doe",
    "mail": "john.doe@example.com"
  },
  ...
  "ldapPassword": {
    "$crypto": {
      "value": {
        "algorithm": "SHA-256",
        "data": "uFacwvr8JsiDwlfI7I/5M+q6jTmQT8e5BaNqXLRcVR+8JxA+/fqy0c8Wo0GhzIz6"
      },
      "type": "salted-hash"
    }
  },
}
}
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--header "Content-Type: application/json" \
--request PATCH \
--data '[ {
  "operation" : "replace",
  "field" : "ldapPassword",
  "value" : "TTestw0rd2"
} ]' \
"http://localhost:8080/openidm/managed/user/jdoe"
{
  "_id": "jdoe",
  "_rev": "4",
  "userName": "jdoe",
  "givenName": "John",
  "sn": "Doe",
  "displayName": "John Doe",
  ...
  "ldapPassword": {
    "$crypto": {
      "value": {
        "algorithm": "SHA-256",
        "data": "kzzz6Nc38srk28xhaBLNX1DDtVsauKnDERoXyVy/TSYtEiMwd2KitgTn7498LZs0"
      },
      "type": "salted-hash"
    }
  }
}
}
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--header "Content-Type: application/json" \
--request PATCH \
--data '[ {
  "operation" : "replace",
  "field" : "ldapPassword",
  "value" : "TTestw0rd3"
} ]' \
"http://localhost:8080/openidm/managed/user/jdoe"
{

```

```

    "_id": "jdoe",
    "_rev": "5",
    "userName": "jdoe",
    "givenName": "John",
    "sn": "Doe",
    "displayName": "John Doe",
    ...
    "ldapPassword": {
      "$crypto": {
        "value": {
          "algorithm": "SHA-256",
          "data": "5NEEkfSsUHF0yEHa/C6yXl9x8s3Q5yaLYJgF02Lp/hPQ8DBKmwUU0U37cqFLQLQX"
        },
        "type": "salted-hash"
      }
    }
  }
}

```

User jdoe now has a *history* of `ldapPassword` values, that contains `TTestw0rd3`, `TTestw0rd2`, `TTestw0rd1`, and `TTestw0rd`, in that order. You can see the four separate hashed values in the `fieldHistory` property of the user:

```

$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/managed/user/jdoe?_fields=fieldHistory"
{
  "_id": "jdoe",
  "_rev": "5",
  "fieldHistory": {
    ...
    "ldapPassword": [
      {
        "$crypto": {
          "value": {
            "algorithm": "SHA-256",
            "data": "k1A1udvQo2gAW/5HxFFjs+IG2p34prv36UsVP89YAVv/bALQTAUJjBhml+qrLLBx"
          },
          "type": "salted-hash"
        }
      },
      {
        "$crypto": {
          "value": {
            "algorithm": "SHA-256",
            "data": "LWHaTZYSUp6hP1RChZElfHmfFBQVv+FGtZuHJxsda/j8s0vjyqeGxk+17IFCX/0l"
          },
          "type": "salted-hash"
        }
      },
      {
        "$crypto": {
          "value": {
            "algorithm": "SHA-256",
            "data": "I4nR+Kkh0s053Sy2V7SUc6Hv3eETC9d0HWopgDTBc9FqRZCV2C9ML0kXGJk8Fhfv"
          },
          "type": "salted-hash"
        }
      }
    ]
  }
}

```

```

    }
  },
  {
    "$crypto": {
      "value": {
        "algorithm": "SHA-256",
        "data": "um9iNdwU7XEVArep2X5I0wr4rRy7nacKXNuzz0c70a1f+lINHkKwZKxaTyBwGbpX2"
      },
      "type": "salted-hash"
    }
  }
]
}
}

```

- The history size for the `ldapPassword` policy is set to 2. To demonstrate the history policy, attempt to patch `jdoe`'s entry with a password value that was used in his previous 2 password resets:

`TTestw0rd2`:

```

$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--header "Content-Type: application/json" \
--request PATCH \
--data '[ {
  "operation" : "replace",
  "field" : "ldapPassword",
  "value" : "TTestw0rd2"
} ]' \
"http://localhost:8080/openidm/managed/user/jdoe"
{
  "code": 403,
  "reason": "Forbidden",
  "message": "Failed policy validation",
  "detail": {
    "result": false,
    "failedPolicyRequirements": [
      {
        "policyRequirements": [
          {
            "policyRequirement": "IS_NEW"
          }
        ],
        "property": "ldapPassword"
      }
    ]
  }
}
}

```

The password reset fails the `IS_NEW` policy requirement.

- Now, reset `jdoe`'s password to a value that was not used in the previous two updates:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--header "Content-Type: application/json" \
--request PATCH \
--data '[ {
  "operation" : "replace",
  "field" : "ldapPassword",
  "value" : "TTestw0rd"
} ]' \
"http://localhost:8080/openidm/managed/user/jdoe"
{
  "_id": "jdoe",
  "_rev": "5",
  "userName": "jdoe",
  "givenName": "John",
  "sn": "Doe",
  "displayName": "John Doe",
  ...
  "ldapPassword": {
    "$crypto": {
      "value": {
        "algorithm": "SHA-256",
        "data": "5NEEkfSsUHF0yEHa/C6yXl9x8s3Q5yaLYJgF02Lp/hPQ8DBKmwUU0U37cqFlQLQX"
      },
      "type": "salted-hash"
    }
  }
}
```

This time, the password reset succeeds.

Chapter 4

Samples That Use the Groovy Connector Toolkit to Create Scripted Connectors

OpenIDM includes a generic Groovy Connector Toolkit that enables you to run Groovy scripts on 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 implementation. Specific scripts are provided within these samples, which demonstrate how the Groovy Connector Toolkit can be used. These scripts cannot be used "as is" in your deployment, but are a good starting point on which to base your customization.

To facilitate creating your own scripted connectors with the Groovy Connector Toolkit, OpenIDM provides a scripted connector *bundler*. The first sample in this chapter uses the connector bundler to create a new connector, and its configuration. The connector bundler is described in detail in the *OpenICF Developers Guide*.

4.1. Using the Connector Bundler to Build a ScriptedSQL Connector

This sample demonstrates the following OpenIDM functionality:

- Custom scripted connector bundler

The sample uses the custom scripted connector bundler to create a new custom connector. The connector bundler generates a scripted connector, the connector configuration and the Groovy scripts required to communicate with an external MySQL database (HRDB).

- Complex data types

Complex data types can be stored, retrieved and synchronized like any other object property. They are stored in the managed data as JSON objects, represented as a string, but can be mapped to external resources in any format required. You can customize the mapping to do additional work with or transformations on the complex data types.

This sample defines one complex data type, *cars*, discussed in more detail later in this section.

- Event hooks to perform actions

The mapping from the internal repository to the external `hrdb` database (`managedUser_systemHrdb`), (defined in the `sync.json` file), includes two script hooks. The first hook is for an `onCreate` event and the second is for an `onUpdate` event. Using these event hooks, OpenIDM logs a statement to the log when a user is created or updated in the external system. In this sample, the script source is included in the mapping. However, a script can also be called from an external file. For more information on event hooks, see Section E.2, "Places to Trigger Scripts" in the *Integrator's Guide*.

- Custom scripted endpoints

All scripted connectors support the configuration of custom scripted endpoints. These are configured in the provisioner configuration file and allow you to execute custom scripts over REST. This example uses a custom scripted endpoint to reset the database and populate it with data. Custom scripted endpoints are illustrated in the custom script step of Section 4.1.2, "Building the Custom ScriptedSQL Connector".

Caution

Because MySQL cannot "un-hash" user passwords there is no way for a reconciliation operation to retrieve and store the password from MySQL and store it in the managed user object. This issue might impact configurations that support multiple external resources in that passwords might not be synchronized immediately after reconciliation from MySQL to the managed/user repository. Users who are missing from managed/user will be created by the reconciliation but their passwords will be empty. When those users are synchronized to other external resources, they will have empty passwords in those resources. Additional scripting might be required to handle this situation, depending on the requirements of your deployment.

The Groovy scripts required for the sample are located in the `sample3/tools` directory. You will need to customize these scripts to address the requirements of your specific deployment, however, the sample scripts are a good starting point on which to base your customization.

The scripted connector bundler takes a configuration file, in JSON format (`sample3/data/scriptedsql.json`). This file includes the details of the connection to the MySQL server, and the list of object types and properties that will be used in the sample. You can use this configuration file as the basis for creating your own connector with the custom scripted connector bundler.

4.1.1. Before You Start

Before you start with this sample, complete the following steps:

- Prepare a fresh installation of OpenIDM. (See Section 1.3, "Preparing the Server").
- Download and install the Apache Maven build tool.
- Configure an external MySQL database, as follows:
 1. Download MySQL Connector/J, version 5.1 or later from the MySQL website. Unpack the delivery, and copy the `.jar` into the `openidm/bundle` directory.

```
$ cp mysql-connector-java-version-bin.jar /path/to/openidm/bundle/
```

2. Set up MySQL to listen on localhost, port 3306. You will connect to the database as user `root` with password `password`.

If you want to use an existing MySQL instance that runs on a different host or port, adjust the configuration file for the sample (`sample3/data/scriptedsql.json`) before you launch the connector bundler. The default generated configuration is as follows:

```
"configurationProperties" : {  
  "username" : "root",  
  "password" : "password",  
  "driverClassName" : "com.mysql.jdbc.Driver",  
  "url" : "jdbc:mysql://localhost:3306/hrdb",
```

3. Create the `hrdb` database, with which OpenIDM will synchronize its managed user repository.

```
$ mysql -u root -p  
Enter password:  
Welcome to the MySQL monitor.  Commands end with ; or \g.  
Your MySQL connection id is 58  
Server version: 5.7.10  
  
Copyright (c) 2000, 2015, Oracle and/or its affiliates. All rights reserved.  
  
Oracle is a registered trademark of Oracle Corporation and/or its  
affiliates. Other names may be trademarks of their respective  
owners.  
  
Type 'help;' or '\h' for help. Type '\c' to clear the current input statement.  
mysql> CREATE DATABASE hrdb CHARACTER SET utf8 COLLATE utf8_bin;  
Query OK, 1 row affected (0.00 sec)  
mysql> quit  
Bye
```

4.1.2. Building the Custom ScriptedSQL Connector

This section uses the custom scripted connector bundler to generate the classes and configuration files required to build a new connector. The custom connector that you build in this section will be used to complete the sample.

1. Create a new directory named `create-connector` in the `openidm/samples/sample3` directory and change to that new directory.

```
$ mkdir /path/to/openidm/samples/sample3/create-connector  
$ cd /path/to/openidm/samples/sample3/create-connector
```

2. Run the custom scripted connector bundler `.jar`, with the configuration file for this sample (`sample3/data/scriptedsql.json`).

```
$ java -jar ../../tools/custom-scripted-connector-bundler-5.0.0.jar -c ../data/scriptedsql.json  
Custom Scripted Connector Bundler for OpenIDM v5.0.0  
Generating connector sources for HRDB-ScriptedSQLConnector
```

This step generates a Maven project (`pom.xml` file) and a `src` directory that contains the packages to be bundled into the connector.

3. In addition to the generated packages, you must add the scripts required to perform operations on your resource. The scripts to access the resource illustrated in this sample are provided in the `sample3/tools` directory. Copy these scripts into the generated `resources/script/hrdb/` directory, so that they can be bundled with the connector.

```
$ cp ../tools/* src/main/resources/script/hrdb/
```

You can customize these scripts before you bundle them, to suit the requirements of your deployment. For more information about writing Groovy scripts to interact with a resource, see the [OpenICF Developer's Guide](#).

4. Use the Maven build tool to build the custom connector, with the configuration and scripts that you provided in the previous steps.

To run this command, you must be in the `create-connector` directory, in which your Maven project (`pom.xml`) is located.

```
$ mvn install
[INFO] Scanning for projects...
Downloading: http://maven.forgerock.org/repo/releases/org/forgerock/openicf/connectors/
connectors-parent/1.5.0.0/connectors-parent-1.5.0.0.pom
Downloaded: http://maven.forgerock.org/repo/releases/org/forgerock/openicf/connectors/
connectors-parent/1.5.0.0/connectors-parent-1.5.0.0.pom (21 KB at 9.2 KB/sec)
[INFO]
[INFO] -----
[INFO] Building 1.4.1.0
[INFO] -----
...
[INFO] Installing org/forgerock/openicf/connectors/hrdb-connector/1.4.1.0/hrdb-connector-1.4.1.0.jar
[INFO] Writing OBR metadata
[INFO] -----
[INFO] BUILD SUCCESS
[INFO] -----
[INFO] Total time: 48.313 s
[INFO] Finished at: 2015-12-10T14:03:02+02:00
[INFO] Final Memory: 37M/320M
[INFO] -----
```

This step generates a connector `.jar` file (`hrdb-connector-1.4.1.0.jar`) in the `target` directory. This connector `.jar` will be used in the rest of this sample.

5. Copy the new connector `.jar` file to the `openidm/connectors` directory, so that it can be picked up by OpenIDM.

```
$ cd /path/to/openidm/samples/sample3
$ cp create-connector/target/hrdb-connector-1.4.1.0.jar ../../connectors/
```

You now have a custom-built connector that includes all the required files for it to be displayed in the OpenIDM Admin UI. The bundled connector also includes the scripts and provisioner configuration that enable it to be used with OpenIDM.

6. Extract the connector configuration file (`provisioner.openicf-hrdb.json`) from the bundled connector into your sample's `conf` directory.

```
$ jar -xvf ../../connectors/hrdb-connector-1.4.1.0.jar conf/provisioner.openicf-hrdb.json
inflated: conf/provisioner.openicf-hrdb.json
```

7. The generated connector configuration file includes no system actions by default.

Edit the value of the `systemActions` property in the connector configuration file, to call a custom script (`tools/ResetDatabaseScript.groovy`) over the REST interface. This script will reset the `hrdb` database and populate it with sample data.

The edited excerpt of the `conf/provisioner.openicf-hrdb.json` file should appear as follows:

```
"systemActions": [
  {
    "scriptId": "ResetDatabase",
    "actions": [
      {
        "systemType": ".*HRDBConnector",
        "actionType": "Groovy",
        "actionFile": "tools\\ResetDatabaseScript.groovy"
      }
    ]
  }
],
```

Currently, only Groovy scripts are supported for these types of actions.

8. Finally, add the generated HTML template file to the UI extensions folder, to enable the new connector to be viewed and configured in the Admin UI.

Inside the connector jar, locate the file that contains the string `1.4.html`.

```
$ cd /path/to/openidm
$ jar -tvf connectors/hrdb-connector-1.4.1.0.jar | grep "1.4.html"
12703 Wed Nov 09 17:42:00 SAST 2016 ui/org.forgerock.openicf.connectors.hrdb.HRDBConnector_1.4.html
```

Create a new extension directory for the connector template.

```
$ mkdir -p ui/admin/extension/templates/admin/connector
```

Extract the HTML template file that you found in the preceding step and then move it into that directory

```
$ jar -xvf connectors/hrdb-connector-1.4.1.0.jar ui/org.forgerock.openicf.connectors.hrdb.HRDBConnector_1.4.html
inflated: ui/org.forgerock.openicf.connectors.hrdb.HRDBConnector_1.4.html
$ mv ui/org.forgerock.openicf.connectors.hrdb.HRDBConnector_1.4.html ui/admin/extension/templates/admin/connector
```

4.1.3. Run the Sample

1. Start OpenIDM with the configuration for sample 3.

```
$ cd /path/to/openidm
$ ./startup.sh -p samples/sample3
Executing ./startup.sh...
Using OPENIDM_HOME: /path/to/openidm
Using PROJECT_HOME: /path/to/openidm/samples/sample3/
Using OPENIDM_OPTS: -Xmx1024m -Xms1024m
Using LOGGING_CONFIG: -Djava.util.logging.config.file=/path/to/openidm/samples/sample3//conf/logging
.properties
Using boot properties at /path/to/openidm/samples/sample3/conf/boot/boot
.properties
-> OpenIDM ready
```

2. Run the custom script described in the previous section to reset the database and populate it with sample data.

You can run the script again, at any point, to reset the database.

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request POST \
"http://localhost:8080/openidm/system/hrdb?_action=script&scriptId=ResetDatabase"
{
  "actions": [
    {
      "result": "Database reset successful."
    }
  ]
}
```

The `hrdb` database should now be populated with sample data.

You can review the contents of the database as follows:

```

$ mysql -u root -p
Enter password:
...
mysql > use hrdb;
Reading table information for completion of table and column names
You can turn off this feature to get a quicker startup with -A

Database changed
mysql > select * from users;
+-----+-----+-----+-----+-----+-----+
| id | uid | password | firstname | lastname | fullname | email |
+-----+-----+-----+-----+-----+-----+-----+
| 1 | bob | e38ad2149... | Bob | Fleming | Bob Fleming | Bob.Fle...
| 2 | rowley | 2aa60a8ff... | Rowley | Birkin | Rowley Birkin | Rowley...
| 3 | louis | 1119cfd37... | Louis | Balfour | Louis Balfour | Louis.B...
| 4 | john | ald7584da... | John | Smith | John Smith | John.Sm...
| 5 | jdoe | edba955d0... | John | Doe | John Doe | John.Do...
+-----+-----+-----+-----+-----+-----+
5 rows in set (0.00 sec)

```

Note

The passwords in the output shown above are hashed to the SHA-1 standard, as they cannot be read into OpenIDM as clear text. The SHA-1 Hash function is used for compatibility reasons. Use a more secure algorithm in a production database.

4.1.4. Reconciling the Repository

1. The mapping configuration file (`sync.json`) for this sample includes the mapping `systemHrdb_managedUser`, which synchronizes users from the source `hrdb` database with the target OpenIDM repository.

You can test this part of the sample by using the `curl` command-line utility, or the OpenIDM Administration UI.

- To reconcile the repository by using the Administration UI:
 1. Log in to the Admin UI at the URL <https://localhost:8443/admin> as the default administrative user (`openidm-admin`) with password `openidm-admin`.

Warning

To protect your deployment in production, change the default administrative password. To do so, select Self-Service from the dropdown list at the top right of the screen and click Change Password.

Return to the Admin View to continue with the sample. (Select Admin View from the top right dropdown list.)

2. Select Configure > Mappings.

The Mappings page shows two configured mappings, one from the **hrdb** database to the OpenIDM repository (**managed/user**), and one in the opposite direction.

3. Click the first mapping (systemHrdb_managedUser) and click Reconcile.

- To reconcile the repository by using the command-line, launch the reconciliation operation with the following command:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request POST \
"http://localhost:8080/openidm/recon?
action=recon&mapping=systemHrdb_managedUser&waitForCompletion=true"
{
  "state": "SUCCESS",
  "_id": "f3c618aa-cc3b-49ed-9a3a-00b012db2513"
}
```

The reconciliation operation creates the five users from the MySQL database in the OpenIDM repository.

2. Retrieve the list of users from the repository.

- To retrieve the users in the repository from the Admin UI:
 1. Select Manage > User to display the User List.

The five users from the **hrdb** database have been reconciled to the OpenIDM repository.

2. To retrieve the details of a specific user, click that user entry.

- To retrieve the users from the repository by using the command-line, query the IDs in the repository as follows:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/managed/user?_queryId=query-all-ids"
{
  "result": [
    {
      "_id": "9d7c304a-fd89-4b58-bd6a-99b2a6a94691",
      "_rev": "1"
    },
    {
      "_id": "53479e98-5460-421c-9e81-0f3a7cc45881",
      "_rev": "1"
    },
    {
      "_id": "4103b904-c7d6-45c2-a9ca-8e563a975fa8",
      "_rev": "1"
    },
    {
      "_id": "1ea17866-aaed-4c51-b3a8-5fa8eb600e04",
      "_rev": "1"
    },
    {
      "_id": "074588a6-64f8-4cce-bb2f-33490aab90ae",
      "_rev": "1"
    }
  ],
  ...
}
```

To retrieve a complete user record, query the `userName` of the individual user entry. The following query returns the record for the user **Rowley Birkin**:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/managed/user/?_queryId=for-username&uid=rowley"
{
  "result": [
    {
      "_id": "53479e98-5460-421c-9e81-0f3a7cc45881",
      "_rev": "1",
      "mail": "Rowley.Birkin@example.com",
      "userName": "rowley",
      "sn": "Birkin",
      "organization": "SALES",
      "givenName": "Rowley",
      "cars": [
        {
          "year": "2013",
          "make": "BMW",
          "model": "328ci"
        },
        {
          "year": "2010",
          "make": "Lexus",
          "model": "ES300"
        }
      ],
      "accountStatus": "active"
    },
    ...
  ]
}
```

Regardless of how you have retrieved Rowley Birkin's entry, note the `cars` property in this user's entry. This property demonstrates a complex object, stored in JSON format in the user entry, as a list that contains multiple objects. In the MySQL database, the `car` table joins to the `users` table through a `cars.users_id` column. The Groovy scripts read this data from MySQL and repackage it in a way that OpenIDM can understand. With support for complex objects, the data is passed through to OpenIDM as a list of `car` objects. Data is synchronized from OpenIDM to MySQL in the same way. Complex objects can also be nested to any depth.

Group membership (not demonstrated here) is maintained with a traditional "join table" in MySQL (`groups_users`). OpenIDM does not maintain group membership in this way, so the Groovy scripts do the work to translate membership between the two resources.

4.1.5. Using Paging With Sample 3

All OpenICF connectors from version 1.4 onwards support the use of paging parameters to restrict query results. The following command indicates that only two records should be returned (`_pageSize=2`) and that the records should be sorted according to their `timestamp` and `_id` (`_sortKeys=timestamp,id`). Including the `timestamp` in the sort ensures that, as you page through the set, changes to records that have already been visited are not lost. Instead, those records are pushed onto the last page:

```
$ curl \
```

```
--header "X-OpenIDM-Username: openidm-admin" \  
--header "X-OpenIDM-Password: openidm-admin" \  
--request GET \  
"http://localhost:8080/openidm/system/hrdb/account?_queryFilter=true&_pageSize=2&_sortKeys=timestamp,id"  
{  
  "result": [  
    {  
      "_id": "1",  
      "email": "Bob.Fleming@example.com",  
      "cars": [  
        {  
          "year": "1979",  
          "make": "Ford",  
          "model": "Pinto"  
        }  
      ],  
      "uid": "bob",  
      "organization": "HR",  
      "firstName": "Bob",  
      "fullName": "Bob Fleming",  
      "lastName": "Fleming"  
    },  
    {  
      "_id": "2",  
      "email": "Rowley.Birkin@example.com",  
      "cars": [  
        {  
          "year": "2013",  
          "make": "BMW",  
          "model": "328ci"  
        }  
      ],  
      "uid": "rowley",  
      "organization": "SALES",  
      "firstName": "Rowley",  
      "fullName": "Rowley Birkin",  
      "lastName": "Birkin"  
    }  
  ],  
  "resultCount": 2,  
  "pagedResultsCookie": "2016-11-09 17:58:50.0,2",  
  "totalPagedResultsPolicy": "NONE",  
  "totalPagedResults": -1,  
  "remainingPagedResults": -1  
}
```

The `pagedResultsCookie` is used by the server to keep track of the position in the search results. You can ignore the `"remainingPagedResults": -1` in the output. The real value of this property is not returned because the scripts that the connector uses do not do any counting of the records in the resource.

Using the `pagedResultsCookie` from the previous step, run a similar query, to retrieve the following set of records in the database. Note that the value of the `pagedResultsCookie` must be URL-encoded, as shown in the following example:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/system/hrdb/account?_queryId=query-all-ids&_pageSize=2&_sortKeys=timestamp,id&_pagedResultsCookie=2016-11-09+17%3A58%3A50.0%2C2"
{
  "result": [
    {
      "_id": "3",
      "uid": "louis"
    },
    {
      "_id": "4",
      "uid": "john"
    }
  ],
  "resultCount": 2,
  "pagedResultsCookie": "2016-11-09 17:58:50.0,4",
  "totalPagedResultsPolicy": "NONE",
  "totalPagedResults": -1,
  "remainingPagedResults": -1
}
```

For more information about paging support, see Section 8.3.5, "Paging and Counting Query Results" in the *Integrator's Guide*.

4.2. Using the Groovy Connector Toolkit to Connect to OpenDJ With ScriptedREST

This sample uses the Groovy Connector Toolkit to implement a ScriptedREST connector, which interacts with the OpenDJ REST API.

The Groovy Connector Toolkit is bundled with OpenIDM in the JAR `openidm/connectors/groovy-connector-1.4.3.0.jar`.

The connector configuration file for this sample (`samples/scriptedrest2dj/conf/provisioner.openicf-scriptedrest.json`) indicates the ScriptedREST implementation of the Groovy connector as follows:

```
{
  "name": "scriptedrest",
  "connectorRef": {
    "connectorHostRef": "#LOCAL",
    "connectorName": "org.forgerock.openicf.connectors.scriptedrest.ScriptedRESTConnector",
    "bundleName": "org.forgerock.openicf.connectors.groovy-connector",
    "bundleVersion": "[1.4.0.0,1.5.0.0)"
  },
  ...
}
```

The Groovy scripts required for the sample are located in the `samples/scriptedrest2dj/tools` directory. You will need to customize these scripts to address the requirements of your specific deployment, however, the sample scripts are a good starting point on which to base your customization.

Important

The Rest2ldap HTTP endpoint provided with OpenDJ is an evolving interface. As such, compatibility between versions is not guaranteed. This sample was tested with OpenDJ 5.

4.2.1. Setting Up OpenDJ

This sample assumes an OpenDJ server, running on the localhost. Follow these steps to install and configure an OpenDJ instance.

1. Download and extract the OpenDJ zip archive from ForgeRock's BackStage site.
2. Set up OpenDJ as follows. Include the `--httpPort` option to enable HTTP access during the setup. This example uses port `8090` so that it does not conflict with the default OpenIDM port:

```
$ cd /path/to/opendj
$ ./setup \
  directory-server \
  --rootUserDN "cn=Directory Manager" \
  --rootUserPassword password \
  --hostname localhost \
  --ldapPort 1389 \
  --adminConnectorPort 4444 \
  --httpPort 8090 \
  --addBaseEntry \
  --baseDN dc=com \
  --acceptLicense
Validating parameters .....Done.
Configuring Certificates .....Done.
Configuring server .....Done.
Creating Base Entry dc=com .....Done.
Starting Directory Server .....Done.
```

To see basic server status and configuration, you can launch `/path/to/opendj/bin/status`

The sample assumes the following OpenDJ configuration:

- The server is installed on the localhost.
 - The server listens for LDAP connections on port 1389.
 - The administration connector port is 4444.
 - The root user DN is `cn=Directory Manager`.
 - The root user password is `password`.
3. Configure the OpenDJ server for replication.

To enable liveSync, this server must be configured for replication, even if it does not actually participate in a replication topology. The following commands configure the server for replication.

```
$ cd /path/to/opendj/bin
$ ./dsconfig create-replication-server \
--hostname localhost \
--port 4444 \
--bindDN "cn=Directory Manager" \
--bindPassword password \
--provider-name "Multimaster Synchronization" \
--set replication-port:8989 \
--set replication-server-id:2 \
--type generic \
--trustAll \
--no-prompt

$ ./dsconfig create-replication-domain \
--hostname localhost \
--port 4444 \
--bindDN "cn=Directory Manager" \
--bindPassword password \
--provider-name "Multimaster Synchronization" \
--domain-name example_com \
--set base-dn:dc=example,dc=com \
--set replication-server:localhost:8989 \
--set server-id:3 \
--type generic \
--trustAll \
--no-prompt
```

4. Enable the OpenDJ HTTP access log.

```
$ ./dsconfig \
set-log-publisher-prop \
--hostname localhost \
--port 4444 \
--bindDN "cn=Directory Manager" \
--bindPassword password \
--publisher-name "File-Based HTTP Access Logger" \
--set enabled:true \
--no-prompt \
--trustAll
```

5. Import the LDIF data required for the sample.

```
$ ./ldapmodify \
--bindDN "cn=Directory Manager" \
--bindPassword password \
--hostname localhost \
--port 1389 \
--filename /path/to/openidm/samples/scriptedrest2dj/data/ldap.ldif
Processing ADD request for dc=example,dc=com
ADD operation successful for DN dc=example,dc=com
Processing ADD request for ou=Administrators,dc=example,dc=com
ADD operation successful for DN ou=Administrators,dc=example,dc=com
Processing ADD request for uid=idm,ou=Administrators,dc=example,dc=com
ADD operation successful for DN uid=idm,ou=Administrators,dc=example,dc=com
Processing ADD request for ou=People,dc=example,dc=com
ADD operation successful for DN ou=People,dc=example,dc=com
Processing ADD request for ou=Groups,dc=example,dc=com
ADD operation successful for DN ou=Groups,dc=example,dc=com
```

6. Set up the access control that enables the OpenIDM administrator user to read the OpenDJ changelog:

```
$ ./dsconfig \
set-access-control-handler-prop \
--hostname localhost \
--port 4444 \
--bindDN "cn=Directory Manager" \
--bindPassword password \
--add global-aci:"(target=\"ldap:///cn=changelog\")(targetattr=\"*|+\" ) \
(version 3.0; acl \"IDM can access cn=changelog\"; \
allow (read,search,compare) \
userdn=\"ldap:///uid=idm,ou=Administrators,dc=example,dc=com\";)" \
--trustAll \
--no-prompt

$ ./dsconfig \
set-access-control-handler-prop \
--hostname localhost \
--port 4444 \
--bindDN "cn=Directory Manager" \
--bindPassword password \
--add global-aci:"(targetcontrol=\"1.3.6.1.4.1.26027.1.5.4\" ) \
(version 3.0; acl \"IDM changelog control access\"; \
allow (read) \
userdn=\"ldap:///uid=idm,ou=Administrators,dc=example,dc=com\";)" \
--trustAll \
--no-prompt
```

7. Enable the default Rest2ldap HTTP endpoint:

```
$ ./dsconfig \  
set-http-endpoint-prop \  
--hostname localhost \  
--port 4444 \  
--bindDN "cn=Directory Manager" \  
--bindPassword password \  
--endpoint-name /api \  
--set authorization-mechanism:"HTTP Basic" \  
--set config-directory:config/rest2ldap/endpoints/api \  
--set enabled:true \  
--no-prompt \  
--trustAll
```

For more information, see *To Set Up REST Access to User Data* in the *OpenDJ Administration Guide*.

8. Replace the default OpenDJ REST to LDAP configuration with the configuration for this sample:

```
$ cd /path/to/openshutdown  
$ cp /path/to/openidm/samples/scriptedrest2dj/data/example-v1.json config/rest2ldap/endpoints/api/
```

9. Restart OpenDJ for the configuration change to take effect.

```
$ cd /path/to/openshutdown/bin  
$ ./stop-ds --restart  
Stopping Server...  
The Directory Server has started successfully  
[07/Nov/2016:14:41:40 +0200] category=PROTOCOL severity=NOTICE msgID=276 msg=Started  
listening for new connections on HTTP Connection Handler 0.0.0.0 port 8090
```

OpenDJ is now configured for this sample.

4.2.2. Running the Sample

This section illustrates the basic CRUD operations on users and groups using the ScriptedREST connector and the OpenDJ REST API. Note that the power of the Groovy connector is in the associated Groovy scripts, and their application in your particular deployment. The scripts provided with this sample are specific to the sample and customization of the scripts is required.

1. Start OpenIDM with the configuration for the ScriptedREST sample:

```
$ cd /path/to/openidm  
$ ./startup.sh -p samples/scriptedrest2dj/
```

2. Check the connector configuration is correct by obtaining the status of the connector, over REST:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request POST \
"http://localhost:8080/openidm/system/scriptedrest?_action=test"
{
  "name": "scriptedrest",
  "enabled": true,
  "config": "config/provisioner.openicf/scriptedrest",
  "objectTypes": [
    "ALL",
    "account",
    "group"
  ],
  "connectorRef": {
    "bundleName": "org.forgerock.openicf.connectors.groovy-connector",
    "connectorName": "org.forgerock.openicf.connectors.scriptedrest.ScriptedRESTConnector",
    "bundleVersion": "[1.4.0.0,1.5.0.0)"
  },
  "displayName": "Scripted REST Connector",
  "ok": true
}
```

3. Create a group entry on the OpenDJ server.

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--header "Content-Type: application/json" \
--request POST \
--data '{
  "cn": "group1"
}' \
"http://localhost:8080/openidm/system/scriptedrest/group?_action=create"
{
  "_id": "group1",
  "cn": "group1",
  "members": null,
  "displayName": "group1",
  "lastModified": null,
  "created": "2017-01-30T11:11:30Z"
}
```

4. Create a user entry on the OpenDJ server. This command creates a user with `uid` scarter:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--header "Content-Type: application/json" \
--request POST \
--data '{
  "givenName" : "Steven",
  "familyName" : "Carter",
  "emailAddress" : "scarter@example.com",
  "telephoneNumber" : "444-444-4444",
  "password" : "Passw0rd",
  "displayName" : "Steven.Carter",
  "uid" : "scarter"
}' \
http://localhost:8080/openidm/system/scriptedrest/account?_action=create
{
  "_id": "scarter",
  "givenName": "Steven",
  "emailAddress": "scarter@example.com",
  "uid": "scarter",
  "groups": null,
  "familyName": "Carter",
  "displayName": "Steven.Carter",
  "created": "2017-01-30T11:12:46Z",
  "telephoneNumber": "444-444-4444"
}
```

Notice that at this stage, the user is not a member of any group.

5. Reconcile the OpenDJ server with the OpenIDM managed user repository:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--header "Content-Type: application/json" \
--request POST \
"http://localhost:8080/openidm/recon?
_action=recon&mapping=systemRestLdapUser_managedUser&waitForCompletion=true"
{
  "_id": "bb69265a-4c57-4f52-8623-db7f206f1351-57",
  "state": "SUCCESS"
}
```

The reconciliation creates a managed user whose `_id` is the same as his OpenDJ `uid`, in this case, `scarter`.

6. Update Steven Carter's managed user entry, by modifying his telephone number:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--header "Content-Type: application/json" \
--request PATCH \
--data '[
  {
    "operation": "replace",
    "field": "telephoneNumber",
    "value": "555-555-5555"
  }
]' \
"http://localhost:8080/openidm/managed/user/scarter"
{
  "_id": "scarter",
  "_rev": "3",
  "userName": "scarter",
  "mail": "scarter@example.com",
  "displayName": "Steven.Carter",
  "telephoneNumber": "555-555-5555",
  "givenName": "Steven",
  "sn": "Carter",
  "accountStatus": "active",
  "effectiveRoles": [],
  "effectiveAssignments": []
}
```

7. The automatic synchronization mechanism between the managed user repository and OpenDJ propagates this change to the OpenDJ server. You can check this change by reading scarter's user entry in OpenDJ:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/system/scriptedrest/account/scarter"
{
  "_id": "scarter",
  "givenName": "Steven",
  "emailAddress": null,
  "uid": "scarter",
  "groups": null,
  "familyName": "Carter",
  "displayName": "Steven.Carter",
  "created": "2017-01-30T11:12:46Z",
  "telephoneNumber": "555-555-5555"
}
```

8. Add Steven Carter to the group you created previously, by updating the group entry:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--header "Content-Type: application/json" \
--header "If-Match: *" \
--request PUT \
--data '{
  "_id": "group1",
  "members": [{"_id": "scarter"}]
}' \
http://localhost:8080/openidm/system/scriptedrest/group/group1
{
  "_id": "group1",
  "cn": "group1",
  "members": [
    {
      "_id": "scarter",
      "displayName": "Steven.Carter"
    }
  ],
  "displayName": "group1",
  "lastModified": "2017-01-30T11:16:00Z",
  "created": "2017-01-30T11:11:30Z"
}
```

9. Read Steven Carter's user entry in OpenDJ, to verify that he is now a member of group1:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
http://localhost:8080/openidm/system/scriptedrest/account/scarter
{
  "_id": "scarter",
  "givenName": "Steven",
  "emailAddress": null,
  "uid": "scarter",
  "groups": [
    {
      "_id": "group1"
    }
  ],
  "familyName": "Carter",
  "displayName": "Steven.Carter",
  "created": "2017-01-30T11:12:46Z",
  "telephoneNumber": "555-555-5555"
}
```

10. Read the group entry to verify its members:


```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/system/scriptedrest/group/group1"
{
  "_id": "group1",
  "cn": "group1",
  "members": [
    {
      "_id": "scarter",
      "displayName": "Steven.Carter"
    }
  ],
  "displayName": "group1",
  "lastModified": "2017-01-30T11:16:00Z",
  "created": "2017-01-30T11:11:30Z"
}
```

11. Reconcile the OpenDJ groups with the OpenIDM managed group repository:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request POST \
"http://localhost:8080/openidm/recon?
action=recon&mapping=systemRestLdapGroup_managedGroup&waitForCompletion=true"
{
  "_id": "08a85706-b191-41bd-89a6-4cd69832af4f-127",
  "state": "SUCCESS"
}
```

12. Read the managed group to verify that the OpenDJ group has been added, and that its members have been reconciled to the managed group repository:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/managed/group/group1"
{
  "_id": "group1",
  "_rev": "1",
  "members": [
    {
      "displayName": "Steven.Carter",
      "_id": "scarter"
    }
  ],
  "displayName": "group1"
}
```

13. Delete the OpenDJ user and group entries, returning the OpenDJ server to its initial state.

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
```

```
--request DELETE \  
http://localhost:8080/oidc/system/scriptedrest/account/scarter  
{  
  "_id": "scarter",  
  "givenName": "Steven",  
  "emailAddress": null,  
  "uid": "scarter",  
  "groups": [  
    {  
      "_id": "group1"  
    }  
  ],  
  "familyName": "Carter",  
  "displayName": "Steven.Carter",  
  "created": "2017-01-30T11:12:46Z",  
  "telephoneNumber": "555-555-5555"  
}  
$ curl \  
--header "X-OpenIDM-Username: oidc-admin" \  
--header "X-OpenIDM-Password: oidc-admin" \  
--request DELETE \  
http://localhost:8080/oidc/system/scriptedrest/group/group1  
{  
  "_id": "group1",  
  "cn": "group1",  
  "members": null,  
  "displayName": "group1",  
  "lastModified": "2017-01-30T11:16:00Z",  
  "created": "2017-01-30T11:11:30Z"  
}
```

4.3. Using the Groovy Connector Toolkit to Connect to OpenDJ With ScriptedCREST

This sample uses the Groovy Connector Toolkit to implement a ScriptedCREST connector, which interacts with the ForgeRock Common REST (CREST) API to connect to an OpenDJ instance. The main difference between a CREST-based API and a generic REST API is that the CREST API is inherently recognizable by all ForgeRock products. As such, the sample can leverage CREST resources in the Groovy scripts, to create CREST requests.

The Groovy Connector Toolkit is bundled with OpenIDM, in the JAR `oidc/connectors/groovy-connector-1.4.3.0.jar`.

The connector configuration file for this sample (`samples/scriptedcrest2dj/conf/provisioner.oidc-scriptedcrest.json`) indicates the ScriptedCREST implementation of the Groovy Connector Toolkit as follows:

```
{
  "name": "scriptedcrest",
  "connectorRef": {
    "connectorHostRef": "#LOCAL",
    "connectorName": "org.forgerock.openicf.connectors.scriptedcrest.ScriptedCRESTConnector",
    "bundleName": "org.forgerock.openicf.connectors.groovy-connector",
    "bundleVersion": "[1.4.0.0,1.5.0.0)"
  },
  ...
}
```

The Groovy scripts required for the sample are located in the `samples/scriptedcrest2dj/tools` directory. You will need to customize these scripts to address the requirements of your specific deployment, however, the sample scripts are a good starting point on which to base your customization.

Important

The Rest2ldap HTTP endpoint provided with OpenDJ is an evolving interface. As such, compatibility between versions is not guaranteed. This sample has been tested with OpenDJ version 5.

4.3.1. Setting Up OpenDJ

This sample assumes an OpenDJ server, running on the localhost. Follow these steps to install and configure an OpenDJ instance.

1. Download and extract the OpenDJ zip archive from ForgeRock's BackStage site.
2. Set up OpenDJ as follows. Include the `--httpPort` option to enable HTTP access during the setup. This example uses port `8090` so that it does not conflict with the default OpenIDM port:

```
$ cd /path/to/opendj
$ ./setup \
  directory-server \
  --rootUserDN "cn=Directory Manager" \
  --rootUserPassword password \
  --hostname localhost \
  --ldapPort 1389 \
  --adminConnectorPort 4444 \
  --httpPort 8090 \
  --addBaseEntry \
  --baseDN dc=com \
  --acceptLicense
Validating parameters .....Done.
Configuring Certificates .....Done.
Configuring server .....Done.
Creating Base Entry dc=com .....Done.
Starting Directory Server .....Done.
```

To see basic server status and configuration, you can launch `/path/to/opendj/bin/status`

The sample assumes the following configuration:

- The server is installed on the localhost.

- The server listens for LDAP connections on port 1389.
 - The administration connector port is 4444.
 - The root user DN is `cn=Directory Manager`.
 - The root user password is `password`.
3. Configure the OpenDJ server for replication.

To enable liveSync, this server must be configured for replication, even if it does not actually participate in a replication topology. The following commands configure the server for replication.

```
$ cd /path/to/openssl/bin
$ ./dsconfig create-replication-server \
  --hostname localhost \
  --port 4444 \
  --bindDN "cn=Directory Manager" \
  --bindPassword password \
  --provider-name "Multimaster Synchronization" \
  --set replication-port:8989 \
  --set replication-server-id:2 \
  --type generic \
  --trustAll \
  --no-prompt

$ ./dsconfig create-replication-domain \
  --hostname localhost \
  --port 4444 \
  --bindDN "cn=Directory Manager" \
  --bindPassword password \
  --provider-name "Multimaster Synchronization" \
  --domain-name example_com \
  --set base-dn:dc=example,dc=com \
  --set replication-server:localhost:8989 \
  --set server-id:3 \
  --type generic \
  --trustAll \
  --no-prompt
```

4. Enable the OpenDJ HTTP access log.

```
$ ./dsconfig \
  set-log-publisher-prop \
  --hostname localhost \
  --port 4444 \
  --bindDN "cn=Directory Manager" \
  --bindPassword password \
  --publisher-name "File-Based HTTP Access Logger" \
  --set enabled:true \
  --no-prompt \
  --trustAll
```

5. Import the LDIF data required for the sample.

```
$ ./ldapmodify \
--bindDN "cn=Directory Manager" \
--bindPassword password \
--hostname localhost \
--port 1389 \
--filename /path/to/openidm/samples/scriptedcrest2dj/data/ldap.ldif
Processing ADD request for dc=example,dc=com
ADD operation successful for DN dc=example,dc=com
Processing ADD request for ou=Administrators,dc=example,dc=com
ADD operation successful for DN ou=Administrators,dc=example,dc=com
Processing ADD request for uid=idm,ou=Administrators,dc=example,dc=com
ADD operation successful for DN uid=idm,ou=Administrators,dc=example,dc=com
Processing ADD request for ou=People,dc=example,dc=com
ADD operation successful for DN ou=People,dc=example,dc=com
Processing ADD request for ou=Groups,dc=example,dc=com
ADD operation successful for DN ou=Groups,dc=example,dc=com
```

6. Set up the access control that enables the OpenIDM administrator user to read the OpenDJ changelog:

```
$ ./dsconfig \
set-access-control-handler-prop \
--hostname localhost \
--port 4444 \
--bindDN "cn=Directory Manager" \
--bindPassword password \
--add global-aci:"(target=\"ldap:///cn=changelog\")(targetattr=\"*|+\" ) \
(version 3.0; acl \"IDM can access cn=changelog\"; \
allow (read,search,compare) \
userdn=\"ldap:///uid=idm,ou=Administrators,dc=example,dc=com\";)" \
--trustAll \
--no-prompt

$ ./dsconfig \
set-access-control-handler-prop \
--hostname localhost \
--port 4444 \
--bindDN "cn=Directory Manager" \
--bindPassword password \
--add global-aci:"(targetcontrol=\"1.3.6.1.4.1.26027.1.5.4\" ) \
(version 3.0; acl \"IDM changelog control access\"; \
allow (read) \
userdn=\"ldap:///uid=idm,ou=Administrators,dc=example,dc=com\";)" \
--trustAll \
--no-prompt
```

7. Enable the default Rest2ldap HTTP endpoint:

```
$ ./dsconfig \  
set-http-endpoint-prop \  
--hostname localhost \  
--port 4444 \  
--bindDN "cn=Directory Manager" \  
--bindPassword password \  
--endpoint-name /api \  
--set authorization-mechanism:"HTTP Basic" \  
--set config-directory:config/rest2ldap/endpoints/api \  
--set enabled:true \  
--no-prompt \  
--trustAll
```

For more information, see *To Set Up REST Access to User Data* in the *OpenDJ Administration Guide*.

8. Replace the default OpenDJ REST to LDAP configuration with the configuration for this sample:

```
$ cd /path/to/opendj  
$ cp /path/to/openidm/samples/scriptedcrest2dj/data/example-v1.json config/rest2ldap/endpoints/api/
```

9. Restart OpenDJ for the configuration change to take effect.

```
$ cd /path/to/opendj/bin  
$ ./stop-ds --restart  
Stopping Server...  
....  
The Directory Server has started successfully  
[30/Jan/2017:13:28:05 +0200] category=PROTOCOL severity=NOTICE msgID=276 msg=Started  
listening for new connections on HTTP Connection Handler 0.0.0.0 port 8090
```

OpenDJ is now configured for this sample.

4.3.2. Running the Sample

This section illustrates the basic CRUD operations on users and groups using the ScriptedCREST connector implementation and the OpenDJ REST API. Note that the power of the Groovy connector is in the associated Groovy scripts, and their application in your specific deployment. The scripts provided with this sample are specific to the sample and customization of the scripts is required.

1. Start OpenIDM with the configuration for the ScriptedCREST sample.

```
$ cd /path/to/openidm  
$ ./startup.sh -p samples/scriptedcrest2dj
```

2. Check the connector configuration is correct by obtaining the status of the connector, over REST.

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request POST \
"http://localhost:8080/openidm/system/scriptedcrest?_action=test"
{
  "name": "scriptedcrest",
  "enabled": true,
  "config": "config/provisioner.openicf/scriptedcrest",
  "objectTypes": [
    "groups",
    "_ALL_",
    "users"
  ],
  "connectorRef": {
    "bundleName": "org.forgerock.openicf.connectors.groovy-connector",
    "connectorName": "org.forgerock.openicf.connectors.scriptedcrest.ScriptedCRESTConnector",
    "bundleVersion": "[1.4.0.0,1.5.0.0)"
  },
  "displayName": "Scripted CREST Connector",
  "ok": true
}
```

3. Create a group entry on the OpenDJ server.

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--header "Content-Type: application/json" \
--request POST \
--data '{
  "_id" : "group1"
}' \
"http://localhost:8080/openidm/system/scriptedcrest/groups?_action=create"
{
  "_id": "group1",
  "_rev": "000000002f333c9f",
  "displayName": "group1"
}
```

4. Create a user entry on the OpenDJ server.

```
$ curl \
--header "Content-Type: application/json" \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request POST \
--data '{
  "name": {
    "familyName": "Carter",
    "givenName" : "Steven"
  },
  "contactInformation": {
    "emailAddress" : "scarter@example.com",
    "telephoneNumber" : "444-444-4444"
  },
  "password" : "TestPassw0rd",
}
```

```
    "displayName" : "Steven.Carter",
    "_id" : "scarter"
  }' \
  "http://localhost:8080/openidm/system/scriptedcrest/users?_action=create"
{
  "_id": "scarter",
  "_rev": "00000000105e85d3",
  "contactInformation": {
    "emailAddress": "scarter@example.com",
    "telephoneNumber": "444-444-4444"
  },
  "userName": "scarter@example.com",
  "name": {
    "givenName": "Steven",
    "familyName": "Carter"
  },
  "displayName": "Steven.Carter"
}
```

Notice that at this stage, the user is not a member of any group.

5. Reconcile the OpenDJ server with the OpenIDM managed user repository:

```
$ curl \
  --header "X-OpenIDM-Username: openidm-admin" \
  --header "X-OpenIDM-Password: openidm-admin" \
  --header "Content-Type: application/json" \
  --request POST \
  "http://localhost:8080/openidm/recon?
  _action=recon&mapping=systemCrestLdapUser_managedUser&waitForCompletion=true"
{
  "_id": "f2fc8502-aeaa-4183-a9f8-2b54662d44b5-88",
  "state": "SUCCESS"
}
```

The reconciliation creates a managed user whose `_id` is the same as his OpenDJ `_id`, in this case, `scarter`.

6. Update Steven Carter's managed user entry, by modifying his telephone number:


```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--header "Content-Type: application/json" \
--request PATCH \
--data '[
  {
    "operation" : "replace",
    "field" : "telephoneNumber",
    "value" : "555-555-5555"
  }
]' \
"http://localhost:8080/openidm/managed/user/scarter"
{
  "_id": "scarter",
  "_rev": "3",
  "userName": "scarter@example.com",
  "mail": "scarter@example.com",
  "displayName": "Steven.Carter",
  "telephoneNumber": "555-555-5555",
  "givenName": "Steven",
  "sn": "Carter",
  "accountStatus": "active",
  "effectiveRoles": [],
  "effectiveAssignments": []
}
```

7. The automatic synchronization mechanism between the managed user repository and OpenDJ propagates this change to the OpenDJ server. You can check this change by reading scarter's user entry in OpenDJ:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/system/scriptedcrest/users/scarter"
{
  "_id": "scarter",
  "_rev": "000000006bf49bb2",
  "contactInformation": {
    "telephoneNumber": "555-555-5555"
  },
  "name": {
    "givenName": "Steven",
    "familyName": "Carter"
  },
  "displayName": "Steven.Carter"
}
```

8. Add Steven Carter to the group you created previously, by updating the members of the group entry.

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--header "Content-Type: application/json" \
--header "If-Match: *" \
--request PUT \
--data '{
  "_id": "group1",
  "members": [{"_id": "scarter"}]
}' \
"http://localhost:8080/openidm/system/scriptedcrest/groups/group1"
{
  "_id": "group1",
  "_rev": "00000000230d6f5b",
  "members": [
    {
      "displayName": "Steven.Carter",
      "_id": "scarter"
    }
  ],
  "displayName": "group1"
}
```

9. Read Steven Carter's entry, to verify that he is now a member of group1.

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/system/scriptedcrest/users/scarter"
{
  "_id": "scarter",
  "_rev": "000000006bf49bb2",
  "contactInformation": {
    "telephoneNumber": "555-555-5555"
  },
  "name": {
    "givenName": "Steven",
    "familyName": "Carter"
  },
  "displayName": "Steven.Carter",
  "groups": [
    {
      "_id": "group1"
    }
  ]
}
```

10. Read the group entry to verify its members.

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/system/scriptedcrest/groups/group1"
{
  "_id": "group1",
  "_rev": "00000000230d6f5b",
  "members": [
    {
      "displayName": "Steven.Carter",
      "_id": "scarter"
    }
  ],
  "displayName": "group1"
}
```

11. Reconcile the OpenDJ groups with the OpenIDM managed group repository:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request POST \
"http://localhost:8080/openidm/recon?
action=recon&mapping=systemCrestLdapGroup_managedGroup&waitForCompletion=true"
{
  "_id": "f2fc8502-aeaa-4183-a9f8-2b54662d44b5-224",
  "state": "SUCCESS"
}
```

12. Read the managed group to verify that the OpenDJ group has been and its members have been reconciled to the managed group repository:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/managed/group/group1"
{
  "_id": "group1",
  "_rev": "1",
  "members": [
    {
      "displayName": "Steven.Carter",
      "_id": "scarter"
    }
  ],
  "displayName": "group1"
}
```

13. Delete the user and group entries, returning the OpenDJ server to its initial state.

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request DELETE \
"http://localhost:8080/openidm/system/scriptedcrest/users/scarter"
{
```

```
"_id": "scarter",
"_rev": "000000006bf49bb2",
"contactInformation": {
  "telephoneNumber": "555-555-5555"
},
"name": {
  "givenName": "Steven",
  "familyName": "Carter"
},
"displayName": "Steven.Carter",
"groups": [
  {
    "_id": "group1"
  }
]
}
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request DELETE \
"http://localhost:8080/openidm/system/scriptedcrest/groups/group1"
{
  "_id": "group1",
  "_rev": "00000000230d6f5b",
  "displayName": "group1"
}
```

Chapter 5

Samples That Use the PowerShell Connector Toolkit to Create Scripted Connectors

OpenICF provides a generic PowerShell Connector Toolkit that enables you to run PowerShell scripts on any external resource. 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. You can use the PowerShell Connector Toolkit to create connectors that can provision any Microsoft system. This chapter describes sample connector implementations that enable you to provision to an Active Directory server instance, and to Azure Active Directory (Azure AD).

The PowerShell Connector Toolkit is available from ForgeRock's BackStage site.

5.1. Connect to Active Directory

This sample provides a number of PowerShell scripts that enable you to perform basic CRUD (create, read, update, delete) operations on an Active Directory server. The samples use the MS Active Directory PowerShell module. For more information on this module, see the corresponding [Microsoft documentation](#).

This sample assumes that OpenIDM is running on a Windows system on the localhost. It also assumes that Active Directory and the OpenICF .NET connector server run on a remote Windows server. The PowerShell connector runs on the .NET connector server.

To use this sample for OpenIDM instances installed on UNIX systems, adjust the relevant commands shown with PowerShell prompts.

5.1.1. Setting Up the PowerShell Active Directory Sample

Run the commands in this procedure from the PowerShell command line. The continuation character used in the command is the back-tick (`).

1. Install, configure, and start the .NET connector server on the machine that is running an Active Directory Domain Controller or on a workstation on which the Microsoft Active Directory PowerShell module is installed.

For instructions on installing the .NET connector server, see Procedure 13.3, "Installing the .NET Connector Server" in the *Integrator's Guide*.

2. Configure OpenIDM to connect to the .NET connector server.

To do so, copy the remote connector provisioner file from the `openidm\samples\provisioners` directory to your project's `conf\` directory, and edit the file according to your configuration.

```
PS C:\> cd \path\to\openidm
PS C:\path\to\openidm> cp samples\provisioners\provisioner.openicf.connectorinfoprovider.json conf
```

For instructions on editing this file, see Procedure 13.7, "Configuring OpenIDM to Connect to the .NET Connector Server" in the *Integrator's Guide*.

3. Download the PowerShell Connector Toolkit archive (`mspowershell-connector-1.4.3.0.zip`) from ForgeRock's BackStage site.

Extract the archive and move the `MsPowerShell.Connector.dll` to the folder in which the connector server application (`connectorserver.exe`) is located.

4. Copy the PowerShell scripts from the `samples\powershell2AD\tools` directory, to the machine on which the connector server is installed.

```
PS C:\path\to\openidm> dir samples\powershell2AD\tools
Directory: C:\path\to\openidm\samples\powershell2AD\tools
```

Mode	LastWriteTime	Length	Name
----	-----	-----	----
----	11/15/2015 01:55 PM	2813	ADAuthenticate.ps1
----	11/15/2015 01:55 PM	10019	ADCreate.ps1
----	11/15/2015 01:55 PM	2530	ADDelete.ps1
----	11/15/2015 01:55 PM	2617	ADResolveUsername.ps1
----	11/15/2015 01:55 PM	7998	ADSchema.ps1
----	11/15/2015 01:55 PM	3706	ADSearch.ps1
----	11/15/2015 01:55 PM	4827	ADSync.ps1
----	11/15/2015 01:55 PM	2075	ADTest.ps1
----	11/15/2015 01:55 PM	10044	ADUpdate.ps1

```
PS C:\path\to\openidm>
```

5. Copy the sample connector configuration for the PowerShell connector from the `samples\provisioners` directory to your project's `conf` directory.

```
PS C:\> cd \path\to\openidm
PS C:\> cp samples\provisioners\provisioner.openicf-adpowershell.json conf
```

The following excerpt of the sample connector configuration shows the configuration properties:

```
"configurationProperties" : {
  "AuthenticateScriptFileName" : "C:/openidm/samples/powershell2AD/tools/ADAuthenticate.ps1",
  "CreateScriptFileName" : "C:/openidm/samples/powershell2AD/tools/ADCreate.ps1",
  "DeleteScriptFileName" : "C:/openidm/samples/powershell2AD/tools/ADDelete.ps1",
  "ResolveUsernameScriptFileName" : "C:/openidm/samples/powershell2AD/tools/ADResolveUsername.ps1",
  "SchemaScriptFileName" : "C:/openidm/samples/powershell2AD/tools/ADSchema.ps1",
  "SearchScriptFileName" : "C:/openidm/samples/powershell2AD/tools/ADSearch.ps1",
  "SyncScriptFileName" : "C:/openidm/samples/powershell2AD/tools/ADSync.ps1",
  "TestScriptFileName" : "C:/openidm/samples/powershell2AD/tools/ADTest.ps1",
  "UpdateScriptFileName" : "C:/openidm/samples/powershell2AD/tools/ADUpdate.ps1",
  "VariablesPrefix" : "Connector",
  "QueryFilterType" : "AdPsModule",
  "ReloadScriptOnExecution" : true,
  "UseInterpretersPool" : true,
  "SubstituteUidAndNameInQueryFilter" : true,
  "UidAttributeName" : "ObjectGUID",
  "NameAttributeName" : "DistinguishedName",
  "PsModulesToImport" : [ "ActiveDirectory" ],
  "Host" : "",
  "Port" : null,
  "Login" : "",
  "Password" : null,
  "CustomProperties" : [ "baseContext = CN=Users,DC=example,DC=com" ],
  "MinInterpretersPoolSize" : 1,
  "MaxInterpretersPoolSize" : 10
},
```

The sample connector configuration assumes that the scripts are located in `C:/openidm/samples/powershell2AD/tools/`. If you copied your scripts to a different location, or are using a different base context for search and synchronization operations such as `DC=example,DC=org`, adjust your connector configuration file accordingly.

Note that the OpenICF framework requires the path to use forward slash characters and not the backslash characters that you would expect in a Windows path.

The host, port, login and password of the machine on which Active Directory runs do not need to be specified here. By default the Active Directory cmdlets pick up the first available Domain Controller. In addition, the scripts are executed using the credentials of the .Net connector server.

Note

The `"ReloadScriptOnExecution"` property is set to `true` in this sample configuration. This setting causes script files to be reloaded each time the script is invoked. Having script files reloaded each time is suitable for debugging purposes. However, this property should be set to `false` in production environments, as the script reloading can have a negative performance impact.

In addition, make sure that the value of the `"connectorHostRef"` property in the connector configuration file matches the value that you specified in the remote connector configuration file, in step 2 of this procedure. For example:

```
"connectorHostRef" : "dotnet",
```

5.1.2. Testing the PowerShell Active Directory Sample

Because you have copied all of the required configuration files into the default OpenIDM project, you can start OpenIDM with the default configuration (that is, without the `-p` option).

```
PS C:\ cd \path\to\openidm
PS C:\ .\startup.bat
```

When OpenIDM has started, you can test the sample by using the `curl` command-line utility. The following examples test the scripts that were provided in the `tools` directory.

1. Test the connector configuration, and whether OpenIDM is able to connect to the .NET connector server with the following request.

```
PS C:\ curl `
--header "X-OpenIDM-Username: openidm-admin" `
--header "X-OpenIDM-Password: openidm-admin" `
--request POST `
"http://localhost:8080/openidm/system?_action=test"
[
  {
    "ok": true,
    "connectorRef": {
      "bundleVersion": "[1.4.2.0,1.5.0.0)",
      "bundleName": "MsPowerShell.Connector",
      "connectorName": "Org.ForgeRock.OpenICF.Connectors.MsPowerShell.MsPowerShellConnector"
    },
    "objectTypes": [
      "_ALL_",
      "group",
      "account"
    ],
    "config": "config/provisioner.openicf/adpowershell",
    "enabled": true,
    "name": "adpowershell"
  }
]
```

2. Query the users in your Active Directory with the following request:

```
PS C:\ curl `
--header "X-OpenIDM-Username: openidm-admin" `
--header "X-OpenIDM-Password: openidm-admin" `
--request GET `
"http://localhost:8080/openidm/system/adpowershell/account?_queryId=query-all-ids"
{
  "remainingPagedResults": -1,
  "pagedResultsCookie": null,
  "resultCount": 1257,
  "result": [
    {
      "_id": "7c41496a-9898-4074-a537-bed696b6be92",
      "distinguishedName": "CN=Administrator,CN=Users,DC=example,DC=com"
    },
    {
      "_id": "f2e08a5c-473f-4798-a2d5-d5cc27c862a9",
      "distinguishedName": "CN=Guest,CN=Users,DC=example,DC=com"
    }
  ]
}
```



```

    },
    {
      "_id": "99de98a3-c125-48dd-a7c2-e21f1488ab06",
      "distinguishedName": "CN=Ben Travis,CN=Users,DC=example,DC=com"
    },
    {
      "_id": "0f7394cc-c66a-404f-ad6d-38dbb4b6526d",
      "distinguishedName": "CN=Barbara Jensen,CN=Users,DC=example,DC=com"
    },
    {
      "_id": "3e6fa858-ed3a-4b58-9325-1fca144eb7c7",
      "distinguishedName": "CN=John Doe,CN=Users,DC=example,DC=com"
    },
    {
      "_id": "6feef4a0-b121-43dc-be68-a96703a49aba",
      "distinguishedName": "CN=Steven Carter,CN=Users,DC=example,DC=com"
    }
  }
  ...

```

- To return the complete record of a specific user, include the ID of the user in the URL. The following request returns the record for Steven Carter.

```

PS C:\ curl `
--header "X-OpenIDM-Username: openidm-admin" `
--header "X-OpenIDM-Password: openidm-admin" `
--request GET `
"http://localhost:8080/openidm/system/adpowershell/account/6feef4a0-b121-43dc-be68-a96703a49aba"
{
  "_id": "6feef4a0-b121-43dc-be68-a96703a49aba",
  "postalCode": null,
  "passwordNotRequired": false,
  "cn": "Steven Carter",
  "name": "Steven Carter",
  "trustedForDelegation": false,
  "uSNChanged": "47219",
  "manager": null,
  "objectGUID": "6feef4a0-b121-43dc-be68-a96703a49aba",
  "modifyTimeStamp": "11/27/2014 3:37:16 PM",
  "employeeNumber": null,
  "sn": "Carter",
  "userAccountControl": 512,
  "passwordNeverExpires": false,
  "displayName": "Steven Carter",
  "initials": null,
  "pwdLastSet": "130615726366949784",
  "scriptPath": null,
  "badPasswordTime": "0",
  "employeeID": null,
  "badPwdCount": "0",
  "accountExpirationDate": null,
  "userPrincipalName": "steve.carter@ad0.example.com",
  "sAMAccountName": "steve.carter",
  "mail": "steven.carter@example.com",
  "logonCount": "0",
  "cannotChangePassword": false,
  "division": null,
  "streetAddress": null,

```

```

"allowReversiblePasswordEncryption": false,
"description": null,
"whenChanged": "11/27/2014 3:37:16 PM",
"title": null,
"lastLogon": "0",
"company": null,
"homeDirectory": null,
"whenCreated": "6/23/2014 2:50:48 PM",
"givenName": "Steven",
"telephoneNumber": "555-2518",
"homeDrive": null,
"uSNCreated": "20912",
"smartcardLogonRequired": false,
"distinguishedName": "CN=Steven Carter,CN=Users,DC=example,DC=com",
"createTimeStamp": "6/23/2014 2:50:48 PM",
"department": null,
"memberOf": [
  "CN=employees,DC=example,DC=com"
],
"homePhone": null
}

```

4. Test whether you can authenticate as one of the users in your Active Directory. The username that you specify here can be either an ObjectGUID, UPN, sAMAccountname or CN.

```

$ PS C:\ curl `
--header "X-OpenIDM-Username: openidm-admin" `
--header "X-OpenIDM-Password: openidm-admin" `
--request POST `
"http://localhost:8080/openidm/system/adpowershell/account?action=authenticate&username=Steven+Carter&password=Passw0rd"
{
  "_id": "6feef4a0-b121-43dc-be68-a96703a49aba"
}

```

The request returns the ObjectGUID if the authentication is successful.

5. You can return the complete record for a specific user, using the query filter syntax described in Section 8.3.4, "Constructing Queries" in the *Integrator's Guide*.

The following query returns the record for the guest user.

```

PS C:\ curl `
--header "X-OpenIDM-Username: openidm-admin" `
--header "X-OpenIDM-Password: openidm-admin" `
--request GET `
"http://localhost:8080/openidm/system/adpowershell/account?_queryFilter=cn+eq+"guest""
{
  "remainingPagedResults": -1,
  "pagedResultsCookie": null,
  "resultCount": 1,
  "result": [
    {
      "_id": "f2e08a5c-473f-4798-a2d5-d5cc27c862a9",
      "postalCode": null,
      "passwordNotRequired": true,
      "cn": "Guest",

```

```
"name": "Guest",
"trustedForDelegation": false,
"uSNChanged": "8197",
"manager": null,
"objectGUID": "f2e08a5c-473f-4798-a2d5-d5cc27c862a9",
"modifyTimeStamp": "6/9/2014 12:35:16 PM",
"employeeNumber": null,
"userAccountControl": 66082,
"whenChanged": "6/9/2014 12:35:16 PM",
"initials": null,
"pwdLastSet": "0",
"scriptPath": null,
"badPasswordTime": "0",
"employeeID": null,
"badPwdCount": "0",
"accountExpirationDate": null,
"sAMAccountName": "Guest",
"logonCount": "0",
"cannotChangePassword": true,
"division": null,
"streetAddress": null,
"allowReversiblePasswordEncryption": false,
"description": "Built-in account for guest access to the computer/domain",
"userPrincipalName": null,
"title": null,
"lastLogon": "0",
"company": null,
"homeDirectory": null,
"whenCreated": "6/9/2014 12:35:16 PM",
"givenName": null,
"homeDrive": null,
"uSNCreated": "8197",
"smartcardLogonRequired": false,
"distinguishedName": "CN=Guest,CN=Users,DC=example,DC=com",
"createTimeStamp": "6/9/2014 12:35:16 PM",
"department": null,
"memberOf": [
  "CN=Guests,CN=Builtin,DC=example,DC=com"
],
"homePhone": null,
"displayName": null,
"passwordNeverExpires": true
}
]
}
```

6. Test whether you are able to create a user on the Active Directory server by sending a POST request with the `create` action.

The following request creates the user `Jane Doe` on the Active Directory server.

```
PS C:\ curl `
--header "X-OpenIDM-Username: openidm-admin" `
--header "X-OpenIDM-Password: openidm-admin" `
--header "Content-Type: application/json" `
--request POST `
--data "{
  \"distinguishedName\" : \"CN=Jane Doe,CN=Users,DC=example,DC=com\",
```

```

{"sn\" : \"Doe\",
 \"cn\" : \"Jane Doe\",
 \"sAMAccountName\" : \"sample\",
 \"userPrincipalName\" : \"janedoe@example.com\",
 \"_ENABLE_\" : true,
 \"_PASSWORD_\" : \"Passw0rd\",
 \"telephoneNumber\" : \"0052-611-091\"
}
\"http://localhost:8080/openidm/system/adpowershell/account?action=create\"
{
  \"_id\": \"42725210-8dce-4fdf-b0e0-393cf0377fdf\",
  \"title\": null,
  \"uSNCreated\": \"47244\",
  \"pwdLastSet\": \"130615892934093041\",
  \"cannotChangePassword\": false,
  \"telephoneNumber\": \"0052-611-091\",
  \"smartcardLogonRequired\": false,
  \"badPwdCount\": \"0\",
  \"department\": null,
  \"distinguishedName\": \"CN=Jane Doe,CN=Users,DC=example,DC=com\",
  \"badPasswordTime\": \"0\",
  \"employeeID\": null,
  \"cn\": \"Jane Doe\",
  \"division\": null,
  \"description\": null,
  \"userPrincipalName\": \"janedoe@example.com\",
  \"passwordNeverExpires\": false,
  \"company\": null,
  \"memberOf\": [],
  \"givenName\": null,
  \"streetAddress\": null,
  \"sn\": \"Doe\",
  \"initials\": null,
  \"logonCount\": \"0\",
  \"homeDirectory\": null,
  \"employeeNumber\": null,
  \"objectGUID\": \"42725210-8dce-4fdf-b0e0-393cf0377fdf\",
  \"manager\": null,
  \"lastLogon\": \"0\",
  \"trustedForDelegation\": false,
  \"scriptPath\": null,
  \"allowReversiblePasswordEncryption\": false,
  \"modifyTimeStamp\": \"11/27/2014 8:14:53 PM\",
  \"whenCreated\": \"11/27/2014 8:14:52 PM\",
  \"whenChanged\": \"11/27/2014 8:14:53 PM\",
  \"accountExpirationDate\": null,
  \"name\": \"Jane Doe\",
  \"displayName\": null,
  \"homeDrive\": null,
  \"passwordNotRequired\": false,
  \"createTimeStamp\": \"11/27/2014 8:14:52 PM\",
  \"uSNChanged\": \"47248\",
  \"sAMAccountName\": \"sample\",
  \"userAccountControl\": 512,
  \"homePhone\": null,
  \"postalCode\": null
}

```

7. Test whether you are able to update a user object on the Active Directory server by sending a PUT request with the complete object, and including the user ID in the URL.

The following request updates user **Jane Doe**'s entry, including her ID in the request. The update sends the same information that was sent in the `create` request, but adds an `employeeNumber`.

```
PS C:\ curl `
--header "X-OpenIDM-Username: openidm-admin" `
--header "X-OpenIDM-Password: openidm-admin" `
--header "Content-Type: application/json" `
--header "If-Match: *" `
--request PUT `
--data "{
  \"distinguishedName\" : \"CN=Jane Doe,CN=Users,DC=example,DC=com\",
  \"sn\" : \"Doe\",
  \"cn\" : \"Jane Doe\",
  \"sAMAccountName\" : \"sample\",
  \"userPrincipalName\" : \"janedoe@example.com\",
  \"_ENABLE_\" : true,
  \"_PASSWORD_\" : \"Passw0rd\",
  \"telephoneNumber\" : \"0052-611-091\",
  \"employeeNumber\" : \"567893\"
}" `
"http://localhost:8080/openidm/system/adpowershell/account/42725210-8dce-4fdf-b0e0-393cf0377fdf"
{
  "_id": "42725210-8dce-4fdf-b0e0-393cf0377fdf",
  "title": null,
  "uSNCreated": "47244",
  "pwdLastSet": "130615906375709689",
  "cannotChangePassword": false,
  "telephoneNumber": "0052-611-091",
  "smartcardLogonRequired": false,
  "badPwdCount": "0",
  "department": null,
  "distinguishedName": "CN=Jane Doe,CN=Users,DC=example,DC=com",
  "badPasswordTime": "0",
  "employeeID": null,
  "cn": "Jane Doe",
  "division": null,
  "description": null,
  "userPrincipalName": "janedoe@example.com",
  "passwordNeverExpires": false,
  "company": null,
  "memberOf": [],
  "givenName": null,
  "streetAddress": null,
  "sn": "Doe",
  "initials": null,
  "logonCount": "0",
  "homeDirectory": null,
  "employeeNumber": "567893",
  "objectGUID": "42725210-8dce-4fdf-b0e0-393cf0377fdf",
  "manager": null,
  "lastLogon": "0",
  "trustedForDelegation": false,
  "scriptPath": null,
  "allowReversiblePasswordEncryption": false,
  "modifyTimeStamp": "11/27/2014 8:37:17 PM",
```

```
"whenCreated": "11/27/2014 8:14:52 PM",
"whenChanged": "11/27/2014 8:37:17 PM",
"accountExpirationDate": null,
"name": "Jane Doe",
"displayName": null,
"homeDrive": null,
"passwordNotRequired": false,
"createTimeStamp": "11/27/2014 8:14:52 PM",
"uSNChanged": "47253",
"sAMAccountName": "sample",
"userAccountControl": 512,
"homePhone": null,
"postalCode": null
}
```

8. Test whether you are able to delete a user object on the Active Directory server by sending a DELETE request with the user ID in the URL.

The following request deletes user *Jane Doe*'s entry.

```
PS C:\ curl `
--header "X-OpenIDM-Username: openidm-admin" `
--header "X-OpenIDM-Password: openidm-admin" `
--request DELETE `
"http://localhost:8080/openidm/system/adpowershell/account/42725210-8dce-4fdf-b0e0-393cf0377fdf"
```

The response includes the complete user object that was deleted.

You can you attempt to query the user object to confirm that it has been deleted.

```
PS C:\ curl `
--header "X-OpenIDM-Username: openidm-admin" `
--header "X-OpenIDM-Password: openidm-admin" `
--request GET `
"http://localhost:8080/openidm/system/adpowershell/account/42725210-8dce-4fdf-b0e0-393cf0377fdf"
{
  "message": "",
  "reason": "Not Found",
  "code": 404
}
```

5.2. Connect to Azure AD

This sample uses the Microsoft Azure Active Directory (Azure AD) PowerShell module. For more information about this module, see <https://msdn.microsoft.com/en-us/library/jj151815.aspx>.

The sample assumes that OpenIDM runs on a local UNIX/Linux machine and that the PowerShell Connector Toolkit (and the OpenICF .NET connector server) run on a remote Windows host with access to an instance of AzureAD. Adjust the command-line examples if your OpenIDM instance runs on Windows.

This sample demonstrates how you can synchronize managed object data such as users and groups with a Microsoft AzureAD deployment.

Note

This sample utilizes a connection between three systems: OpenIDM on UNIX/Linux, an OpenICF .NET connector server on Windows, and Azure AD "in the cloud". Internet connection times vary widely and performance is based on responses to external calls, including potential timeouts, if a command doesn't perform the first time, try again. OpenIDM's synchronization and reconciliation performance will be fully dependent on the performance of the managed resource.

5.2.1. Before You Start

Before you can run this sample, you need to meet several prerequisites. This section describes each of the prerequisites, and how to install or test them.

- You must have a Microsoft account, which gives you access to Microsoft Azure.

You can set up a Microsoft account at <https://signup.live.com/>.

With a Microsoft account, you can access the Azure portal at <http://azure.microsoft.com>.

- You must have an Azure AD cloud directory.

If you do not have an existing Azure AD cloud, set one up as follows:

1. Navigate to <https://account.windowsazure.com/signup>. Once you log in with your Microsoft credentials, fill in the prompts and Microsoft will create an Azure subscription.
2. Navigate to <http://portal.azure.com>, log in with your Microsoft account.
3. In the Microsoft Azure screen, select New on the left hand menu.
4. From the New list, select Security + Identity > Active Directory.
5. Complete the Add Directory form with the details of your directory, and select the check mark at the bottom of the form to submit.

✕

Add directory

DIRECTORY ?

Create new directory ▼

NAME ?

Example

DOMAIN NAME ?

example .onmicrosoft.com

COUNTRY OR REGION ?

United States ▼

This is a B2C directory. ? **PREVIEW**

✓

Your directory should now be created and listed.

- Apart from your default Microsoft Azure account, you must have an *administrative user account* for your Azure AD.

By default your directory will have a single identity, your Microsoft Azure account. You cannot use this account to run the PowerShell Connector scripts that administer the Azure AD.

If your Azure AD does not already include other administrative accounts, create a local administrative identity that is native to your directory as follows:

1. Log in to <https://portal.azure.com/> with your Microsoft Azure credentials.
2. From the left-hand menu, select Browse > Active Directory.
3. Select your cloud directory from the left-hand menu and select USERS in the top navigation bar.

- At the bottom of the page select Add User and enter the details of the new administrative user.

ADD USER

Tell us about this user

TYPE OF USER

New user in your organization

USER NAME ?

admin @ example.onmicrosoft.com

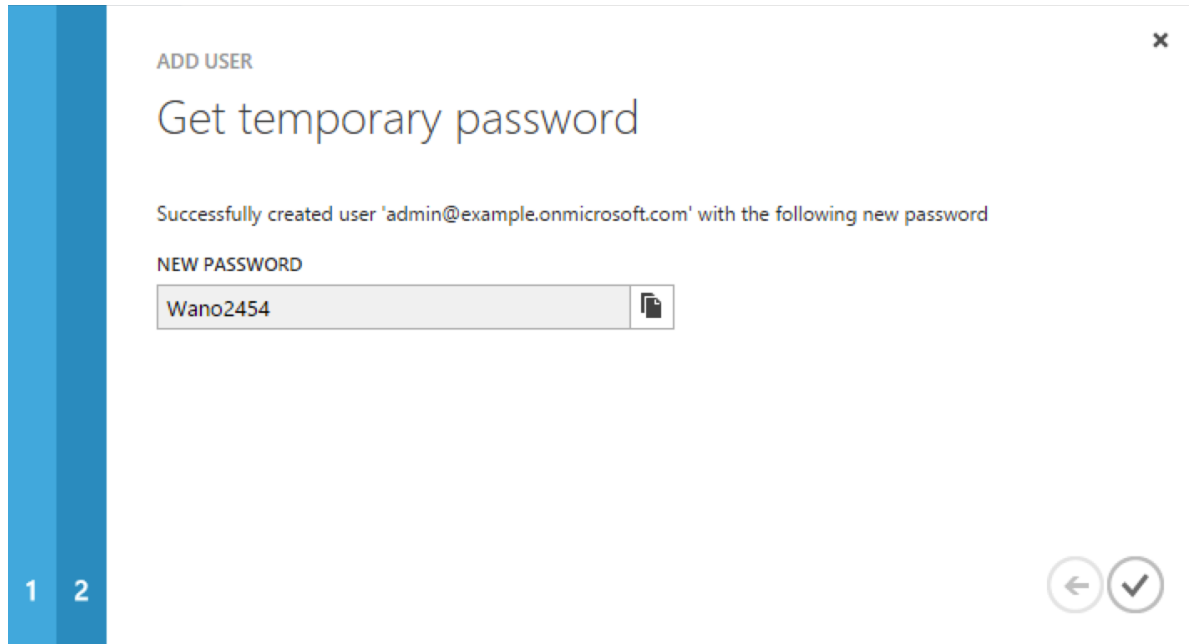
2 3

Select the arrow to continue.

- On the User Profile screen, enter the details of this administrative user. Make sure that the user's Role is at least User Admin.

Select the arrow to continue.

- On the final screen, select Create and note the temporary password that is assigned to the user.



Because new administrative users are forced to change their password on first login, you should log in as this user to change the password.

Select the check mark to complete the new user creation process.

7. Select the username at the top right of the screen and select Sign-out to sign out of your Microsoft Azure account, then select SIGN IN > Use Another Account to sign in as your new administrative user.
8. Enter the email address of the new administrative user and select Continue.
9. Enter the temporary password that you received and select Sign In.
10. On the Update Your Password screen, enter a new password, then select Update password and sign in.

You now have a new administrative user account that the PowerShell scripts will use to access your Azure AD.

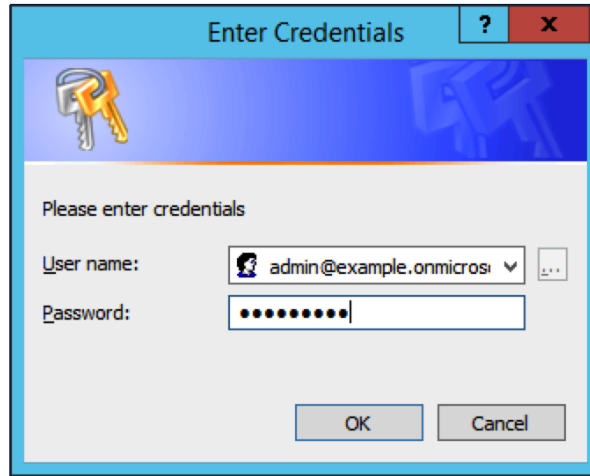
- The Windows Azure AD Module for Windows PowerShell must be installed on the Windows host that connects to Azure.

If needed, install the Azure AD Module as described in the following Microsoft article.

- Your Windows host must be able to contact your Azure AD deployment.

Verify the connection as follows:

1. Open a PowerShell window and type `Connect-MsolService` at the command prompt.
2. On the Enter Credentials screen, enter the credentials of the administrative account that you created for the Azure directory.



If the PowerShell command returns with no error, you have successfully connected to your remote Azure AD deployment.

- The OpenICF .NET connector server must be installed on your Windows host.

If you have not yet installed the .NET connector server, follow the instructions in Section 13.4.1, "Installing and Configuring a .NET Connector Server" in the *Integrator's Guide*.

- The PowerShell Connector Toolkit must be installed on your Windows host.

If you have not yet installed the PowerShell Connector Toolkit, follow the instructions in Chapter 5, "*PowerShell Connector Toolkit*" in the *Connectors Guide*. In these instructions, you will use a command with a `/setkey` option to create a password key for your .NET connector server. You will use that key in Section 5.2.2, "Setting Up the PowerShell Azure AD Sample on OpenIDM".

Important

Before you continue, check that the OpenICF .NET connector server is still running. If it is not running, restart the connector server and check the logs. In some cases, Windows blocks the PowerShell connector dll. If the connector server fails to start, right-click on `MsPowerShell.Connector.dll` and select Properties > Security. If you see the following text on that tab:

This file came from another computer and might be blocked to help protect this computer.

Select the Unblock button to unblock the connector dll. Then restart the connector server.

When all of the above elements are in place, you can proceed with running the sample, as described in Section 5.2.2, "Setting Up the PowerShell Azure AD Sample on OpenIDM".

5.2.2. Setting Up the PowerShell Azure AD Sample on OpenIDM

This section assumes that OpenIDM is installed on the local UNIX/Linux machine.

1. On the Windows host, create a directory for the PowerShell scripts.

The sample connector configuration expects the scripts in the directory `C:/openidm/samples/powershell2AzureAD/tools/`. If you put them in a different location, adjust your connector configuration accordingly.

```
PS C:\> mkdir -Path openidm\samples\powershell2AzureAD\azureADScripts

Directory: C:\openidm\samples\powershell2AzureAD

Mode                LastWriteTime         Length Name
----                -
d-----          5/4/2016  11:26 AM             azureADScripts

PS C:\>
```

2. Copy the PowerShell sample scripts from the OpenIDM instance on your UNIX/Linux host to the new directory on the remote Windows server.

One way to do this is to run an `scp` client, such as `pscp` in your Windows terminal. The following command copies the PowerShell scripts from the OpenIDM installation to the Windows machine:

```
PS C:\> cd openidm\samples\powershell2AzureAD\tools
PS C:\> pscp -r username@openidm-host:path/to/openidm/samples/powershell2AzureAD/azureADScripts/*.ps .
```

The following scripts should now be in the `azureADScripts` directory on your Windows system:

```
PS C:\openidm\samples\powershell2AzureAD\azureADScripts> ls

Directory: C:\openidm\samples\powershell2AzureAD\azureADScripts

Mode                LastWriteTime         Length Name
----                -
-a---              5/4/2016  11:26 AM         7258 AzureADCreate.ps1
-a---              5/4/2016  11:26 AM         3208 AzureADDelete.ps1
-a---              5/4/2016  11:26 AM         6952 AzureADSchema.ps1
-a---              5/4/2016  11:26 AM         8149 AzureADSearch.ps1
-a---              5/4/2016  11:26 AM         2465 AzureADTest.ps1
-a---              5/4/2016  11:26 AM        10840 AzureADUpdate.ps1
```

Note

You need to set the execution policy, as Windows by default does not trust downloaded scripts. For more information, see the following article: [Using the Set-ExecutionPolicy Cmdlet](#)

You can then run the [Unblock-File](#) cmdlet to allow OpenIDM to run the scripts on your Windows system. For more information, see the following article: [Unblock-File](#).

3. On the Linux/UNIX machine on which OpenIDM is installed, navigate to the `path/to/openidm/samples/powershell2AzureAD` directory, and open the `provisioner.openicf.connectorinfoprovider.json conf` file.
4. Edit the remote connector server configuration file to match the settings of the remote .NET connector server.

Change the port to `8760`, and the password (`key`) that you configured for the .NET connector server.

The following example assumes that the .NET connector server is running on the host `198.51.100.1`, listening on the default port, and configured with a secret key of `Passw0rd`:

```
{
  "remoteConnectorServers" :
  [
    {
      "name" : "dotnet",
      "host" : "198.51.100.1",
      "port" : 8760,
      "useSSL" : false,
      "timeout" : 0,
      "key" : "Passw0rd"
    }
  ]
}
```

5. Open the sample Azure AD PowerShell connector configuration file, `provisioner.openidcf-azureadpowershell.json`, and edit it to match your deployment. In particular, set the following properties in that file:

```
"Host" : "198.51.100.1",  
"Port" : 8760,  
"Login" : "admin@example.onmicrosoft.com",  
"Password" : "Passw0rd",
```

Host

The hostname or IP address on which the .NET connector server is running.

Port

The port on which the .NET connector server is listening.

Login

The username of the administrative account you created for the Azure directory in the previous section.

Password

The password of the administrative account you created for the Azure directory in the previous section.

If you have placed the PowerShell scripts in a directory other than the default (`C:\openidm\samples\powershell2AzureAD\azureADScripts`) you must also update those paths in the PowerShell connector configuration file.

6. Start OpenIDM with the PowerShell AzureAD sample configuration:

```
$ cd path/to/openidm  
$ ./startup.sh -p samples/powershell2AzureAD
```

5.2.3. Managing Users and Groups with the PowerShell Azure AD Sample

This section walks you through several REST commands that enable you to test the connector configuration, and perform basic CRUD operations in your Azure AD, through the PowerShell connector.

1. Test that the connector has been configured correctly and that the Azure AD resource can be reached:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request POST \
"http://localhost:8080/openidm/system/azureadpowershell?action=test"
{
  "name": "azureadpowershell",
  "enabled": true,
  "config": "config/provisioner.openicf/azureadpowershell",
  "objectTypes": [
    "ALL",
    "account",
    "group"
  ],
  "connectorRef": {
    "bundleName": "MsPowerShell.Connector",
    "connectorName": "Org.ForgeRock.OpenICF.Connectors.MsPowerShell.MsPowerShellConnector",
    "bundleVersion": "[1.4.2.0,1.5.0.0)"
  },
  "displayName": "PowerShell Connector ",
  "ok": true
}
```

If you see no response from this connector test, review any changes that you made to the `provisioner-openicf*` files in your project's `conf/` subdirectory. If you've made changes appropriate for your deployment, wait a couple of minutes and try again.

2. Query the IDs of the existing users in your Azure AD deployment:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/system/azureadpowershell/account?_queryId=query-all-ids"
{
  "result": [ {
    "_id": "51560d42-e60e-49a8-855b-42b6eca35ca6",
    "UserPrincipalName": "admin@example.onmicrosoft.com"
  },
  {
    "_id": "5e63b42f-c93a-466f-af86-f0a8d00f2491",
    "UserPrincipalName": "scarter@example.onmicrosoft.com"
  } ]
  ,
  ...
}
```

3. Use a query filter to return all details of all existing users in your Azure AD:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/system/azureadpowershell/account?_queryFilter=true"
{
  "result": [
    {
```

```

    "_id": "51560d42-e60e-49a8-855b-42b6eca35ca6",
    "LiveId": "10033FFF96C5186D",
    "FirstName": "Barbara",
    "LastName": "Jensen",
    "UserPrincipalName": "admin@example.onmicrosoft.com",
    "AlternateEmailAddress": [ "bjensen@example.com" ],
    "LastPasswordChangeTimestamp": "3/15/2016 11:02:19 AM",
    "DisplayName": "Barbara Jensen",
    "PasswordNeverExpires": false,
    "MobilePhone": "+1 3602297105"
  },
  {
    "_id": "5e63b42f-c93a-466f-af86-f0a8d00f2491",
    "LiveId": "1003BFFD96A4CFBA",
    "FirstName": "Sam",
    "LastName": "Carter",
    "UserPrincipalName": "scarter@example.onmicrosoft.com",
    "AlternateEmailAddresses": [ "scarter@example.com" ],
    "LastPasswordChangeTimestamp": "3/7/2016 1:09:31 PM",
    "DisplayName": "Sam Carter",
    "PasswordNeverExpires": false,
    "MobilePhone": "+1 3602297105"
  }
]
...}

```

- Return details for a specific user account, by its `_id`

```

$ curl \
  --header "X-OpenIDM-Username: openidm-admin" \
  --header "X-OpenIDM-Password: openidm-admin" \
  --request GET \
  "http://localhost:8080/openidm/system/azureadpowershell/account/51560d42-e60e-49a8-855b-42b6eca35ca6"

```

- Create a new user in Azure AD. Substitute the domain for your Azure AD deployment for `example.onmicrosoft.com`:


```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request POST \
--header "content-type: application/json" \
--data '{
  "PasswordNeverExpires": false,
  "AlternateEmailAddresses": ["John.Bull@example.com"],
  "LastName": "Bull",
  "PreferredLanguage": "en-GB",
  "FirstName": "John",
  "UserPrincipalName": "Dev_John.Bull@example.onmicrosoft.com",
  "DisplayName": "John Bull"
}' \
"http://localhost:8080/openidm/system/azureadpowershell/account?_action=create"
{
  "_id" : "d4aac947-2037-4f29-b0f5-d404fd99938c",
  "LiveId" : "10037FFE979FB2C1",
  "FirstName" : "John",
  "LastName" : "Bull",
  "UserPrincipalName" : "Dev_John.Bull@example.onmicrosoft.com",
  "AlternateEmailAddresses" : [ "John.Bull@example.com" ],
  "LastPasswordChangeTimestamp" : "5/5/2016 3:52:43 PM",
  "DisplayName" : "John Bull",
  "PasswordNeverExpires" : false,
  "PreferredLanguage" : "en-GB"
}
```

Rerun the same command. You should see the following error:

```
{
  "code" : 500,
  "reason" : "Internal Server Error",
  "message" : "Operation CREATE failed with ConnectorException on system object:
    Dev_John.Bull@example.onmicrosoft.com"
}
```

- Update the user entry that you have just created with a patch request. Include the `_id` of the new user in the URL. Save that `_id` value for a later step.

The following example updates the user's display name:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--header "if-match: *" \
--header "content-type: application/json" \
--request PATCH \
--data '[
  {
    "operation": "replace",
    "field": "DisplayName",
    "value": "John P. Bull"
  }
]' \
"http://localhost:8080/openidm/system/azureadpowershell/account/d4aac947-2037-4f29-b0f5-d404fd99938c"
{
  "_id" : "d4aac947-2037-4f29-b0f5-d404fd99938c",
  "LiveId" : "10037FFE979FB2C1",
  "FirstName" : "John",
  "LastName" : "Bull",
  "UserPrincipalName" : "Dev_John.Bull@mikejangfr.onmicrosoft.com",
  "AlternateEmailAddresses" : [ "John.Bull@example.com" ],
  "LastPasswordChangeTimestamp" : "5/5/2016 3:52:43 PM",
  "DisplayName" : "John P. Bull",
  "PasswordNeverExpires" : false,
  "PreferredLanguage" : "en-GB"
}
```

7. Now create a group:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--header 'content-type: application/json' \
--request POST \
--data '{
  "DisplayName" : "Dev Testers group",
  "Description" : "Description of a Dev Group"
}' \
'http://localhost:8080/openidm/system/azureadpowershell/group?action=create'
{
  "_id" : "9091be74-f37e-408d-9198-2d2b5f4b4cdd",
  "Members" : [ ],
  "DisplayName" : "Dev Testers Group",
  "GroupType" : "Security",
  "Description" : "Description of a Dev Group",
  "objectId" : "9091be74-f37e-408d-9198-2d2b5f4b4cdd"
}
```

8. Add your recently created user to this new group. Use the `_id` of that user, as the `ObjectId`. Use the `_id` of the newly created group in the endpoint:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin"
\
--header "X-OpenIDM-Password: openidm-admin"
\
--header "Content-Type: application/json"
\
--header "If-Match: *"
\
--request PUT
\
--data '{
  "Members" : [
    {
      "ObjectId" : "d4aac947-2037-4f29-b0f5-d404fd99938c"
    }
  ]
}' \
"http://localhost:8080/openidm/system/azureadpowershell/group/9091be74-f37e-408d-9198-2d2b5f4b4cdd"
{
  "_id" : "9091be74-f37e-408d-9198-2d2b5f4b4cdd",
  "Members" : [ {
    "ObjectId" : "d4aac947-2037-4f29-b0f5-d404fd99938c",
    "DisplayName" : "John P. Bull",
    "GroupMemberType" : "User",
    "EmailAddress" : "Dev_John.Bull@example.onmicrosoft.com"
  } ],
  "DisplayName" : "Testing Devs Group",
  "GroupType" : "Security",
  "Description" : "Description of a Dev Group",
  "objectId" : "9091be74-f37e-408d-9198-2d2b5f4b4cdd"
}
```

9. Confirm the result, by the `_id` of the group:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin"
\
--header "X-OpenIDM-Password: openidm-admin"
\
--request GET \
"http://localhost:8080/openidm/system/azureadpowershell/group/9091be74-f37e-408d-9198-2d2b5f4b4cdd"
```

10. Update a label for the group. Use the same group `_id`:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--header "Content-Type: application/json" \
--header "If-Match: *" \
--request PUT \
--data '{
  "_id" : "9091be74-f37e-408d-9198-2d2b5f4b4cdd",
  "Description" : "Dev Masters Group",
  "Members" : [
    {
      "ObjectId" : "d4aac947-2037-4f29-b0f5-d404fd99938c",
      "DisplayName" : "John P. Bull",
      "GroupMemberType" : "User",
      "EmailAddress" : "Dev_John.Bull@example.onmicrosoft.com"
    }
  ],
  "DisplayName" : "Testing Devs Group",
  "GroupType" : "Security",
  "objectId" : "9091be74-f37e-408d-9198-2d2b5f4b4cdd"
}' \
"http://localhost:8080/openidm/system/azureadpowershell/group/9091be74-f37e-408d-9198-2d2b5f4b4cdd"
```

You should see the new `Description` in the output.

11. Remove the user from the new group. Use the same group `_id` Note how the `Members` in the `--data` block, and the output, are blank:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin"
\
--header "X-OpenIDM-Password: openidm-admin"
\
--header "Content-Type: application/json"
\
--header "If-Match: *"
\
--request PUT
\
--data '{
  "_id" : "9091be74-f37e-408d-9198-2d2b5f4b4cdd",
  "Description" : "Dev Masters Group",
  "Members" : [ ],
  "DisplayName" : "Testing Devs Group",
  "GroupType" : "Security",
  "objectId" : "9091be74-f37e-408d-9198-2d2b5f4b4cdd"
}' \
"http://localhost:8080/openidm/system/azureadpowershell/group/9091be74-f37e-408d-9198-2d2b5f4b4cdd"
{
  "_id" : "9091be74-f37e-408d-9198-2d2b5f4b4cdd",
  "Members" : [ ],
  "DisplayName" : "Testing Devs Group",
  "GroupType" : "Security",
  "Description" : "Dev Masters Group",
  "objectId" : "9091be74-f37e-408d-9198-2d2b5f4b4cdd"
}
```

12. Delete the user that you created earlier:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin"
\
--header "X-OpenIDM-Password: openidm-admin"
\
--request DELETE \
"http://localhost:8080/openidm/system/azureadpowershell/account/d4aac947-2037-4f29-b0f5-d404fd99938c"
```

To verify that the user was deleted, run the REST call to [query-all-ids](#) shown earlier in this section. The ID associated with that user should have been removed.

5.2.4. Reconciling Users Between OpenIDM and Azure AD

In this section, you'll run commands that demonstrate reconciliation mappings between OpenIDM managed users and your remote instance of Azure AD.

In preparation, create a new user on the Azure AD system:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin"
\
--header "X-OpenIDM-Password: openidm-admin"
\
--request POST
\
--header "content-type: application/json"
\
--data '{
  "UserPrincipalName": "CEO@example.onmicrosoft.com",
  "LastName": "Officer",
  "FirstName": "Chief",
  "DisplayName": "Chief Executive Officer",
  "PasswordNeverExpires": false
}' \
"http://localhost:8080/openidm/system/azureadpowershell/account?_action=create"
```

In the steps that follow, you'll run reconciliations to see what happens to that user in the OpenIDM data store.

1. Review the list of current managed users in the OpenIDM repository, filtered for the `userName` that starts with (sw) CEO:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin"
\
--header "X-OpenIDM-Password: openidm-admin"
\
--request GET \
"http://localhost:8080/openidm/managed/user?_queryFilter=userName+sw+'CEO'"
```

Until you reconcile the Azure AD repository to OpenIDM, the output should be empty:

```
{
  "result" : [ ],
  "resultCount" : 0,
  "pagedResultsCookie" : null,
  "totalPagedResultsPolicy" : "NONE",
  "totalPagedResults" : -1,
  "remainingPagedResults" : -1
}
```

2. Run a reconciliation from Azure AD to OpenIDM:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin"
\
--header "X-OpenIDM-Password: openidm-admin"
\
--request POST \
"http://localhost:8080/openidm/recon?
_action=recon&mapping=systemAzreadpowershellAccount_managedUser&waitForCompletion=true"
{
  "_id" : "71811f1c-2ec0-47ae-ba47-d62c7094201b-1105",
  "state" : "SUCCESS"
}
```

- Now rerun the command to list of current managed users in the OpenIDM repository, filtered for the `userName` that starts with (`sw`) CEO:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin"
\
--header "X-OpenIDM-Password: openidm-admin"
\
--request GET \
"http://localhost:8080/openidm/managed/user?_queryFilter=userName+sw+'CEO'"
{
  "result" : [ {
    "_id" : "3a012a60-19c2-4fb4-99cc-0bb82dc4588c",
    "_rev" : "1",
    "userName" : "CEO@example.onmicrosoft.com",
    "mail" : "CEO@example.onmicrosoft.com",
    "sn" : "Officer",
    "givenName" : "Chief",
    "accountStatus" : "active",
    "effectiveRoles" : [ ],
    "effectiveAssignments" : [ ]
  } ],
  "resultCount" : 1,
  "pagedResultsCookie" : null,
  "totalPagedResultsPolicy" : "NONE",
  "totalPagedResults" : -1,
  "remainingPagedResults" : -1
}
```

- Delete that CEO user from the Azure AD system, by the `_id` shown earlier when you searched that system:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin"
\
--header "X-OpenIDM-Password: openidm-admin"
\
--request DELETE \
"http://localhost:8080/openidm/system/azreadpowershell/account/3a012a60-19c2-4fb4-99cc-0bb82dc4588c"
```

If successful, you'll see the data for the CEO user one last time.

- Run a second reconciliation from the remote Azure AD repository to OpenIDM:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin"
\
--header "X-OpenIDM-Password: openidm-admin"
\
--request POST \
"http://localhost:8080/openidm/recon?
_action=recon&mapping=systemAzureadpowershellAccount_managedUser&waitForCompletion=true"
```

6. Rerun the command to search the OpenIDM repository for a `userName` that starts with 'CEO' one more time, to confirm that user has been reconciled out of the OpenIDM repository:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin"
\
--header "X-OpenIDM-Password: openidm-admin"
\
--request GET \
"http://localhost:8080/openidm/managed/user?_queryFilter=userName+sw+'CEO'"
```


Chapter 6

Audit Samples

This chapter demonstrates two ways to configure OpenIDM to save audit data. In Section 6.1, "Directing Audit Information To a MySQL Database", you'll see how to direct information for OpenIDM audit events to a MySQL database server. In Section 6.2, "Show Audit Events Published on a JMS Topic", you'll see how to configure an audit event handler for the Java Message Service (JMS).

6.1. Directing Audit Information To a MySQL Database

This sample demonstrates how you can use the audit features provided with OpenIDM and directs audit information to a MySQL database. The sample is closely related to Section 2.1, "Reconciling an XML File Resource" but includes a ScriptedSQL implementation of the Groovy Connector Toolkit to connect to the MySQL database.

To use any of the audit features that are demonstrated in this sample, copy information from files in the `samples/audit-sample` directory. This can help you to set up auditing for any other sample.

6.1.1. Audit Sample Configuration Files

Review the configuration files used in this sample. They can help you understand the functionality of the data sets being audited.

The key configuration files, in the `samples/audit-sample` directory, are as follows:

- `conf/provisioner.openicf-scriptedsql.json` shows the configuration of the Scripted SQL implementation of the Groovy Connector. For more information, see Chapter 6, "Groovy Connector Toolkit" in the *Connectors Guide*.
- `conf/provisioner.openicf-xml.json` shows the configuration of the Chapter 15, "XML File Connector" in the *Connectors Guide*.
- `conf/audit.json` configures audit logging on the router, to a remote system, as discussed in Section 21.1, "Configuring the Audit Service" in the *Integrator's Guide*.
- `conf/sync.json` shows mappings between managed users and the data set attached through the XML File Connector.
- `data/sample_audit_db.mysql` includes a schema that supports tables in the external MySQL database.

- Groovy scripts in the `tools/` subdirectory supports communications between the Scripted SQL connector and the MySQL database.

6.1.2. Configuration with MySQL

You need to set up communications between OpenIDM and an external MySQL database server.

Make sure MySQL is running. Follow the configuration instructions described in Section 2.2, "Setting Up a MySQL Repository" in the *Installation Guide*.

The sample expects the following configuration for MySQL:

- The database is available on the local system.
- The database listens on the standard MySQL port, 3306.
- You can connect to the MySQL database over the network.
- OpenIDM should include the MySQL connector .jar file in the `/path/to/openidm/bundle` directory. If you need to download this file, see Section 2.2, "Setting Up a MySQL Repository" in the *Installation Guide* for instructions.
- MySQL serves a database called `audit` with six tables: `auditaccess`, `auditactivity`, `auditauthentication`, `auditconfig`, `auditrecon`, and `auditsync`.
- For more information on the database schema, examine the following data definition language file: `openidm/samples/audit-sample/data/sample_audit_db.mysql`. Import the file into MySQL before running the sample.

```
$ mysql -u root -p < /path/to/openidm/samples/audit-sample/data/sample_audit_db.mysql
Enter password:
$
```

Now review the format of the audit database sample, created from the `sample_audit_db.mysql` file, at the MySQL prompt. To access that prompt, run the following command:

```
$ mysql -u root -p
mysql > use audit;
Reading table information for completion of table and column names
You can turn off this feature to get a quicker startup with -A

Database changed
```

You can now review the current state of each noted audit log with the following commands:

```
select * from auditaccess;
select * from auditactivity;
select * from auditauthentication;
select * from auditconfig;
select * from auditrecon;
select * from auditsync;
```

Unless you enable scheduled reconciliation, you will not see audit reconciliation data until you run reconciliation manually.

6.1.3. Install the Sample

Prepare OpenIDM as described in Section 1.3, "Preparing the Server", then start OpenIDM with the configuration for the audit sample.

```
$ cd /path/to/openidm
$ ./startup.sh -p samples/audit-sample
```

6.1.4. Running the Sample

Run reconciliation over the REST interface.

```
$ curl \
  --header "X-OpenIDM-Username: openidm-admin" \
  --header "X-OpenIDM-Password: openidm-admin" \
  --request POST \
  "http://localhost:8080/openidm/recon?
  _action=recon&mapping=systemXmlfileAccounts_managedUser&waitForCompletion=true"
```

You can also run this reconciliation from the Admin UI, at <https://localhost:8443/admin>.

Warning

The default password for the OpenIDM administrative user, `openidm-admin`, is `openidm-admin`.

To protect your deployment in production, change the default administrative password. To do so, navigate to the Self-Service UI at <https://localhost:8443/> and click Change Password.

You can now review the results in the MySQL database, from the MySQL prompt, using the commands described earlier.

If you have retained the default `"useForQueries" : true` option in the `conf/audit.json` file, you can also `GET` the same data with a REST call. For examples on how you can query audit logs, see Section 21.9, "Querying Audit Logs Over REST" in the *Integrator's Guide*.

6.1.5. Additional Audit Configuration

You can configure a variety of audit event handlers, event topics, audit filter policies, and scripts. You can even set up auditing, by topic, in CSV files. For more information, see Chapter 21, "Logging Audit Information" in the *Integrator's Guide*.

You can see how this works from the Admin UI. After you log in with the OpenIDM administrative account, click Configure > System Preferences > Audit.

6.2. Show Audit Events Published on a JMS Topic

Starting with OpenIDM 4.5.0, you can configure a Java Message Service (JMS) Audit Event Handler. JMS is an API that supports Java-based peer-to-peer messages between clients. The JMS API can create, send, receive, and read messages, reliably and asynchronously. You can now set up a JMS audit event handler, which can publish messages that comply with the *Java(TM) Message Service Specification Final Release 1.1*.

Warning

JMS topics are not related to ForgeRock audit event topics. The ForgeRock implementation of JMS topics use the `publish/subscribe` messaging domain, and may direct messages to the JMS audit event handler. In contrast, ForgeRock audit event topics specify categories of events.

In this sample, we demonstrate the use of the JMS audit event handler. This sample is based on Section 2.1, "Reconciling an XML File Resource". You will set up communications between OpenIDM and an external JMS Message Broker, as well as *Apache Active MQ* as the JMS provider and message broker.

6.2.1. Adding Required Bundles for the JMS Messaging

To test JMS messaging, you'll download the binary distribution of ActiveMQ, as well as four OSGi Bundle JAR files.

- The ActiveMQ 5.13.2 binary, which you can download from <http://activemq.apache.org/activemq-5132-release.html>.
- *The ActiveMQ Client*, 5.13.2 JAR, which you can download from <https://repository.apache.org/content/repositories/releases/org/apache/activemq/activemq-client/>.
- The *bnd* JAR for working with OSGi bundles, which you can download from `bnd-1.50.0.jar`.
- The Apache Geronimo J2EE management bundle, `geronimo-j2ee-management_1.1_spec-1.0.1.jar`, which you can download from http://central.maven.org/maven2/org/apache/geronimo/specs/geronimo-j2ee-management_1.1_spec/1.0.1/.
- The *hawtbuf* Maven-based protocol buffer compiler JAR, which you can download from `hawtbuf-1.11.jar`.

Note

The JMS event handler has been tested and documented with the noted versions of the files that you've just downloaded.

Unpack the `apache-activemq-5.13.2-*` binary.

Make sure at least the first two JAR files, for *the Active MQ Client* and *bnd*, are in the same directory. Navigate to that directory, and create an OSGi bundle with the following steps:

1. Create a BND file named `activemq.bnd` with the following contents:

```
version=5.13.2
Export-Package: *;version=${version}
Bundle-Name: ActiveMQ :: Client
Bundle-SymbolicName: org.apache.activemq
Bundle-Version: ${version}
```

2. Run the following command to create the OSGi bundle archive file:

```
$ java \
-jar \
bnd-1.50.0.jar \
wrap \
-properties \
activemq.bnd \
activemq-client-5.13.2.jar
```

3. Rename the `activemq-client-5.13.2.bar` file that appears to `activemq-client-5.13.2-osgi.jar` and copy it to the `/path/to/openidm/bundle` directory.

Copy the other two bundle files, *Apache Geronimo* and *hawtbuf*, to the `/path/to/openidm/bundle` directory.

6.2.2. Starting the ActiveMQ Broker and OpenIDM

With the appropriate bundles in the `/path/to/openidm/bundles` directory, you're ready to start the ActiveMQ message broker, as well as OpenIDM with the JMS Audit Sample.

Navigate to the directory where you unpacked the ActiveMQ binary, possibly `/path/to/apache-activemq-5.13.0/`. If you need SSL protection for your audit data, edit the ActiveMQ configuration file, `activemq.xml`, in the `conf/` subdirectory. Find the code block associated with `<transportConnectors>`, and add the following line within that block:

```
<transportConnector name="ssl"
uri="ssl://0.0.0.0:61617?transport.enabledCipherSuites=
SSL_RSA_WITH_RC4_128_SHA,SSL_DH_anon_WITH_3DES_EDE_CBC_SHA
&maximumConnections=1000&wireFormat.maxFrameSize=104857600&transport.daemon=true"/>
```

To start the ActiveMQ broker, navigate to the directory where you unpacked the ActiveMQ binary, and run the following command:

```
$ bin/activemq start
INFO: Loading '/path/to/apache-activemq-5.13.0/bin/env'
INFO: Using java '/usr/bin/java'
INFO: Starting - inspect logfiles specified in logging.properties and log4j.properties to get details
INFO: pidfile created : '/path/to/apache-activemq-5.13.0/data/activemq.pid' (pid '22671')
```

Now start OpenIDM, with the sample in the `/path/to/openidm/samples/audit-jms-sample` directory:

```
$ cd /path/to/openidm
$ ./startup.sh -p samples/audit-jms-sample
```

Note

If you see the following error in the OpenIDM console, you may have forgotten to go through the steps shown in Section 6.2.1, "Adding Required Bundles for the JMS Messaging"; you also need to start the ActiveMQ broker.

```
SEVERE: Unable to create JmsAuditEventHandler 'jms': null
```

6.2.3. Configuring and Using a JMS Consumer Application

To take advantage of the Apache ActiveMQ event broker, the JMS audit sample includes a Java consumer in the following directory: `/path/to/openidm/samples/audit-jms-sample/consumer/`

Assuming you have Apache Maven installed on the local system, you can compile that sample consumer with the following commands:

```
$ cd /path/to/openidm/samples/audit-jms-sample/consumer/  
$ mvn clean install
```

When the build process is complete, you'll see a **BUILD SUCCESS** message:

```
[INFO] -----  
[INFO] BUILD SUCCESS  
[INFO] -----  
[INFO] Total time: 12.638 s  
[INFO] Finished at: 2016-04-15T15:18:31-07:00  
[INFO] Final Memory: 13M/119M  
[INFO] -----
```

Note

You may see **[WARNING]** messages during the build. As long as the messages end with **BUILD SUCCESS**, you can proceed with the JMS consumer application.

You can then run the following command to output audit messages related to OpenIDM actions:

```
$ mvn \  
exec:java \  
-Dexec.mainClass="SimpleConsumer"  
\  
-Dexec.args="tcp://localhost:61616"  
[INFO] -----  
[INFO] Building SimpleConsumer 1.0-SNAPSHOT  
[INFO] -----  
[INFO]  
[INFO] --- exec-maven-plugin:1.4.0:java (default-cli) @ SimpleConsumer ---  
Connection factory=org.apache.activemq.ActiveMQConnectionFactory  
READY, listening for messages. (Press 'Enter' to exit)
```

If you've configured ActiveMQ on a secure port, as described in Section 6.2.2, "Starting the ActiveMQ Broker and OpenIDM", you can run this alternative command:

```
$ mvn \
exec:java \
-Dexec.mainClass="SimpleConsumer" \
-Dexec.args="ssl://localhost:61617?daemon=true&socket.enabledCipherSuites=
SSL_RSA_WITH_RC4_128_SHA,SSL_DH_anon_WITH_3DES_EDE_CBC_SHA"
```

Try some actions on OpenIDM, either in a different console or in the Admin UI. Watch the output in the `SimpleConsumer` console. As an example, you might see output similar to the following when you are running reconciliation on the data in this sample:

```
{
  "event": {
    "_id": "88b3da4d-e427-4f21-881c-036d7a854ccc-2559",
    "reconId": "88b3da4d-e427-4f21-881c-036d7a854ccc-2546",
    "mapping": "systemXmlfileAccounts_managedUser",
    "linkQualifier": "default",
    "exception": null,
    "action": "UPDATE",
    "userId": "openidm-admin",
    "eventName": "recon",
    "timestamp": "2016-04-16T13:40:35.974Z",
    "transactionId": "88b3da4d-e427-4f21-881c-036d7a854ccc-2546",
    "message": null,
    "situation": "CONFIRMED",
    "sourceObjectId": "system/xmlfile/account/scarter",
    "status": "SUCCESS",
    "targetObjectId": "managed/user/scarter",
    "reconciling": "source",
    "ambiguousTargetObjectIds": "",
    "entryType": "entry"
  },
  "auditTopic": "recon"
}
```

Chapter 7

Scripted JMS Sample

With OpenIDM, you can subscribe to JMS messaging protocols through its messaging service. When a publisher sends a message over JMS, that message is broadcast. All active and subscribed clients will receive that message.

In an event-driven architecture, also known as a message-driven architecture, there are publishers and subscribers. OpenIDM can publish JMS messages using the JMS Audit Event Handler. OpenIDM can also subscribe to JMS messages using the Messaging Service's JMS Subscriber.

JMS is short for Java Messaging Service. With the scripted message handler configured in this sample, you can configure scripts to parse the contents of JMS messages, and act on that content.

The script in this sample, `crudpaqTextMessageHandler.js`, demonstrates how JMS can handle ForgeRock REST operations. If you customize a script to manage JMS messages, be sure to modify the `messaging.json` file in the `conf/` subdirectory.

This sample uses ActiveMQ, a JMS message broker. With the ActiveMQ UI, you can act as the JMS message provider. This sample demonstrates how you can input REST payloads through the ActiveMQ UI.

To set up ActiveMQ, you must download and process several files. For more information, see Section 6.2.1, "Adding Required Bundles for the JMS Messaging".

7.1. Starting IDM and the ActiveMQ Broker

With the appropriate bundles in the `/path/to/openidm/bundles` directory, you're ready to start the ActiveMQ message broker, as well as OpenIDM with the JMS Audit Sample.

To start the ActiveMQ broker, navigate to the directory where you unpacked the ActiveMQ binary, and run the following command:

```
$ bin/activemq start
INFO: Loading '/path/to/apache-activemq-5.13.0/bin/env'
INFO: Using java '/usr/bin/java'
INFO: Starting - inspect logfiles specified in logging.properties and log4j.properties to get details
INFO: pidfile created : '/path/to/apache-activemq-5.13.0/data/activemq.pid' (pid '22671')
```

Now start OpenIDM, with the sample in the `/path/to/openidm/samples/scriptedJMSSubscriber` directory:

```
$ cd /path/to/openidm
$ ./startup.sh -p samples/scriptedJMSSubscriber
```


7.1.1. Configuring a Secure Port for JMS Messages

To set up SSL protection for your JMS messages, stop ActiveMQ and edit the following files:

- `activemq.xml` in the `/path/to/apache-activemq-5.13.0/` directory, typically where you unpacked the downloaded ActiveMQ binary.

Find the code block associated with `<transportConnectors>`, and add the following line within that block:

```
<transportConnector name="ssl"
  uri="ssl://0.0.0.0:61617?transport.enabledCipherSuites=
  SSL_RSA_WITH_RC4_128_SHA,SSL_DH_anon_WITH_3DES_EDE_CBC_SHA
  &maximumConnections=1000&wireFormat.maxFrameSize=104857600&transport.daemon=true"/>
```

- `messaging.json` in the `conf/` subdirectory of the sample.

Change the value of the `java.naming.provider.url` to:

```
ssl://127.0.0.1:61617?daemon=true&socket.enabledCipherSuites=
SSL_RSA_WITH_RC4_128_SHA,SSL_DH_anon_WITH_3DES_EDE_CBC_SHA
```

After making these changes, start ActiveMQ as described in Section 7.1, "Starting IDM and the ActiveMQ Broker".

7.2. Access the REST Interface via the ActiveMQ UI

In this section, you will run REST calls through the ActiveMQ UI. This assumes you started ActiveMQ and OpenIDM in Section 7.1, "Starting IDM and the ActiveMQ Broker".

1. Access the ActiveMQ web console, at <http://localhost:8161/admin>. Log in as an administrator; the default administrative user and password are `admin` and `admin`.
2. In the ActiveMQ web console, in the Queue Name text box, enter the value of `idmQ` from the `messaging.json` file in the `conf/` subdirectory. The default value for `queue.idmQ` and `destinationName` is `idmQ`. Enter that value and select Create.
3. In the Queues window for the `idmQ` queue, select the `Send To` link under Operations.
4. You should see a `Send a JMS Message` window, with a `Message body` text box towards the bottom.
5. Copy the following text, a payload to create a new user with a user ID of `mgr1`, to the `Message body` text box:

```
{
  "operation" : "CREATE",
  "resourceName" : "/managed/user",
  "newResourceId" : "mgr1",
  "content" : {
    "mail" : "mgr1@example.com",
    "sn" : "Sanchez",
    "givenName" : "Jane",
    "password" : "Password1",
    "employeenumber" : 100,
    "accountStatus" : "active",
    "telephoneNumber" : "",
    "roles" : [ ],
    "postalAddress" : "",
    "userName" : "mgr1",
    "stateProvince" : ""
  },
  "params" : {},
  "fields" : [ "*", "*_ref" ]
}
```

For comparison, note the following equivalent REST call to create the same user:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--header "Content-Type: application/json" \
--request POST \
--data '{
  "mail" : "mgr1@example.com",
  "sn" : "Sanchez",
  "givenName" : "Jane",
  "password" : "Password1",
  "employeenumber" : 100,
  "accountStatus" : "active",
  "telephoneNumber" : "",
  "roles" : [ ],
  "postalAddress" : "",
  "userName" : "mgr1",
  "stateProvince" : "",
  "params" : {},
  "fields" : [ "*", "*_ref" ]
}' \
"http://localhost:8080/openidm/managed/user?_action=create"
```

6. Observe the OpenIDM console. You should see output in two parts. The first part starts with:

```
*****request received*****
Parsed JMS JSON =
...
```

The first part should include a JSON-formatted replay of the request that you entered in the `Message body` text box.

The second part of the output includes information for that same user, as written to the internal OpenIDM managed user datastore.

7. Confirm that user either in the Admin UI, or via the following REST call:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/managed/user/mgr1"
```

8. Repeat the process to create additional users. Try a different REST function. For example, you can enter the following payload in the ActiveMQ UI `Message body` text box, to change the first name (`givenName`) of the `mgr1` user to Donna:

```
{
  "operation" : "PATCH",
  "resourceName" : "/managed/user/mgr1",
  "rev" : "",
  "value" : [
    {
      "operation": "replace",
      "field": "/givenName",
      "value": "Donna"
    }
  ]
}
```

7.3. Customizing the Scripted JMS Sample

If you set up a custom script to parse and process JMS messages, store that script in the `script/` subdirectory. Assume that script is named `myCustomScript.js`.

Edit the `messaging.json` file in the `conf/` subdirectory, and point it to that file:

```

{
  "subscribers" : [
    {
      "name" : "IDM CREST Queue Subscriber",
      "instanceCount": 3,
      "enabled" : true,
      "type" : "JMS",
      "handler" : {
        "type" : "SCRIPTED",
        "properties" : {
          "script" : {
            "type" : "text/javascript",
            "file" : "myCustomScript.js"
          }
        }
      },
      "properties" : {
        "sessionMode" : "CLIENT",
        "jndi" : {
          "contextProperties" : {
            "java.naming.factory.initial" : "org.apache.activemq.jndi.ActiveMQInitialContextFactory",
            "java.naming.provider.url" : "tcp://127.0.0.1:61616?daemon=true",
            "queue.idmQ" : "idmQ"
          },
          "destinationName" : "idmQ",
          "connectionFactoryName" : "ConnectionFactory"
        }
      }
    }
  ]
}

```

You'll find some of these properties in Table J.11, "JMS Audit Event Handler Unique `config` Properties" in the *Integrator's Guide*. Despite the name of the table and the different configuration file, the properties are the same.

Other properties of interest in the `messaging.json` file are shown in the following table:

Table 7.1. JMS `messaging.json` Configuration Properties

<code>messaging.json</code> File Label	Description
<code>subscribers</code>	Needed to subscribe to incoming JMS message requests
<code>name</code>	Arbitrary name for the subscriber
<code>instanceCount</code>	Each <code>instanceCount</code> manages a single connection between OpenIDM and the messaging channel. Supports multithreading throughput. If subscribing to a queue, such as <code>queue.idmQ</code> , the message is handled by a single instance. If subscribing to a topic, all instances receive and handle the same message.
<code>handler</code>	Parses the JMS message, then processes it, possibly through a script

messaging.json File Label	Description
queue.idmQ	One of the JNDI context properties. Name of the JMS queue in the ActiveMQ UI.
destinationName	JNDI lookup name for message delivery

Chapter 8

Demonstrating the Roles Implementation

This chapter illustrates how roles are managed in OpenIDM, and how they can be used to provision to external systems, based on certain criteria.

OpenIDM supports two types of roles:

- *Provisioning roles* - used to specify how objects are provisioned to an external system.
- *Authorization roles* - used to specify the authorization rights of a managed object internally, within OpenIDM.

Both provisioning roles and authorization roles use the relationships mechanism to link the role object, and the managed object to which the role applies. For more information about relationships between objects, see Section 9.5, "Managing Relationships Between Objects" in the *Integrator's Guide*.

There are three samples in this chapter:

- The *provisioning* sample (in `openidm/samples/roles/provrole`) demonstrates the operations that can be performed over the REST interface, or by using the Admin UI, to manage roles in OpenIDM.
- The *crudops* sample (in `openidm/samples/roles/crudops`) demonstrates how attributes are provisioned to an external system (an LDAP directory), based on role membership.
- The *temporalConstraints* sample (in `openidm/samples/roles/temporalConstraints`) builds on the other two roles samples and demonstrates how to apply time-based restrictions to role grants.

The separate samples are described in the sections that follow.

Important

Most of the commands in this sample can be run by using the command-line but it is generally much easier to use the Admin UI. In some cases, the command-line version makes it easier to explain what is happening in OpenIDM. Therefore, in all steps, the sample first shows the command-line version, and then provides the equivalent method of running the command in the Admin UI.

8.1. CRUDOPS Role Sample - Working With Managed Roles

This sample demonstrates how to manage roles by using the REST interface, or the Admin UI. The sample covers the tasks discussed in the following sections:

- Section 8.1.2, "Creating a Managed Role"
- Section 8.1.3, "Reading and Searching Managed Roles"
- Section 8.1.4, "Granting a Managed Role to a User"
- Section 8.1.5, "Removing a Managed User's Roles"
- Section 8.1.6, "Deleting a Managed Role"

8.1.1. Before You Start

This sample does not include its own configuration. Before you work through the sample, start OpenIDM with the default configuration, as follows:

```
$ cd /path/to/openidm
$ ./startup.sh
Executing ./startup.sh...
Using OPENIDM_HOME: /path/to/openidm
Using PROJECT_HOME: /path/to/openidm
Using OPENIDM_OPTS: -Xmx1024m -Xms1024m
Using LOGGING_CONFIG: -Djava.util.logging.config.file=/path/to/openidm/conf/logging.properties
Using boot properties at /path/to/openidm/conf/boot/boot
.properties
-> OpenIDM ready
```

8.1.2. Creating a Managed Role

In this section, you will create two managed roles - an *Employee* role and a *Contractor* role. The sample shows how to create the first role directly over the REST interface and the second by using the Admin UI. Choose whichever method is easiest for you to create both roles.

To create the Employee role by using the REST interface:

```
$ curl \
--header "Content-type: application/json" \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request POST \
--data '{
  "name" : "Employee",
  "description": "Role granted to workers on the payroll."
}' \
"http://localhost:8080/openidm/managed/role?_action=create"
{
  "_id": "ad19979e-adbb-4d35-8320-6db50646b432",
  "_rev": "1",
  "name": "Employee",
  "description": "Role granted to workers on the payroll."
}
```

Take note of the role's identifier (ad19979e-adbb-4d35-8320-6db50646b432 in this example). Although you can create roles with a PUT request and specify your own ID, you should use system-generated IDs in a production environment to avoid conflicts and referential integrity issues.

To create the Contractor role by using the Admin UI:

1. Point your browser to the Admin UI URL (for example <https://localhost:8443/admin/>).
2. On the Dashboard, click Manage Role > New Role.
3. Enter a name for the role (Contractor) and a description.
4. Click Save.

The managed role has been created, but currently has no assignments. These will be added in the second roles sample.

Note

The Admin UI creates managed objects with an immutable, system-generated identifier. You will see this in the next section, in which you read or query the role object.

8.1.3. Reading and Searching Managed Roles

This section shows how to read, and search managed role objects.

The easiest way to see a list of managed role objects is by using the Admin UI:

1. Navigate to the Admin UI URL.
2. On the Dashboard, click Manage Role.

The list of roles (currently only Contractor and Employee) is displayed.

If you know the identifier, it is easy to read a managed role object over the REST interface. The following command reads the Employee role, that you created in the previous section:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/managed/role/ad19979e-adbb-4d35-8320-6db50646b432"
{
  "_id": "ad19979e-adbb-4d35-8320-6db50646b432",
  "_rev": "1",
  "name": "Employee",
  "description": "Role granted to workers on the payroll."
}
```

You can also *search* for the role object, with a filtered query, if you know the role name. For more information about filtered queries, see Section 8.3.1, "Common Filter Expressions" in the *Integrator's Guide*.

The following query retrieves all roles with a name equal to "Contractor". The `_prettyPrint=true` parameter displays the result in a format that is easy to read:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/managed/role?_queryFilter=/name+eq+'Contractor'&_prettyPrint=true"

{
  "result": [
    {
      "_id": "b02d2531-5066-415e-bc90-31fe57e02322",
      "_rev": "1",
      "name": "Contractor",
      "description": "Role granted to contract workers."
    }
  ],
  ...
}
```

To retrieve a list of all the managed roles that have been defined, query the `managed/role` endpoint, as follows:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/managed/role?_queryFilter=true&_prettyPrint=true"

{
  "result" : [ {
    "_id" : "ad19979e-adbb-4d35-8320-6db50646b432",
    "_rev" : "1",
    "name" : "Employee",
    "description" : "Role granted to workers on the payroll."
  }, {
    "_id" : "b02d2531-5066-415e-bc90-31fe57e02322",
    "_rev" : "1",
    "name" : "Contractor",
    "description" : "Role granted to contract workers."
  } ],
  ...
}
```

8.1.4. Granting a Managed Role to a User

For a role to be useful, it must be granted to a managed user. This section creates a new managed user entry, Felicitas Doe, then assigns the Employee role, created in the previous section, to Felicitas's user entry.

1. Create the user entry over REST, as follows:

```
$ curl \
--header "Content-type: application/json" \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request POST \
--data '{
  "mail": "fdoe@example.com",
  "sn": "Doe",
  "telephoneNumber": "555-1234",
  "userName": "fdoe",
  "givenName": "Felicitas",
  "description": "Felicitas Doe",
  "displayName": "fdoe"
}' \
"http://localhost:8080/openidm/managed/user?_action=create"
{
  "_id": "837085ae-766e-417c-9b7e-c36eee4352a3",
  "_rev": "1",
  "mail": "fdoe@example.com",
  "sn": "Doe",
  "telephoneNumber": "555-1234",
  "userName": "fdoe",
  "givenName": "Felicitas",
  "description": "Felicitas Doe",
  "displayName": "fdoe",
  "accountStatus": "active",
  "effectiveRoles": [],
  "effectiveAssignments": []
}
```

Note that Felicitas has no `effectiveRoles` or `effectiveAssignments` by default.

Tip

Create the user entry in the Admin UI:

- Click Manage User > New User from the Dashboard.

2. Grant the Employee role to Felicitas's entry by sending a PATCH request to update her entry, and providing a pointer to the role ID:

```
$ curl \
--header "Content-type: application/json" \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request PATCH \
--data '[
  {
    "operation" : "add",
    "field" : "/roles/-",
    "value" : {"_ref": "managed/role/ad19979e-adbb-4d35-8320-6db50646b432"}
  }
]' \
"http://localhost:8080/openidm/managed/user/837085ae-766e-417c-9b7e-c36eee4352a3"
{
  "_id": "837085ae-766e-417c-9b7e-c36eee4352a3",
  "_rev": "2",
  "mail": "fdoe@example.com",
  "sn": "Doe",
  "telephoneNumber": "555-1234",
  "userName": "fdoe",
  "givenName": "Felicitas",
  "description": "Felicitas Doe",
  "displayName": "fdoe",
  "accountStatus": "active",
  "effectiveRoles": [
    {
      "_ref": "managed/role/ad19979e-adbb-4d35-8320-6db50646b432"
    }
  ],
  "effectiveAssignments": []
}
```

Tip

Grant the role in the Admin UI:

1. Select Manage > User and click on fdoe's entry.
2. Click Provisioning Roles, select the Employee role and click Add Role.

OR

1. Select Manage > Role and click on the Employee Role.
2. Select the Role Members tab, click Add Role Members, browse for fdoe's entry, and click Add.

3. Now, query Felicitas's entry to return her `roles` and `effectiveRoles`.

The `effectiveRoles` property is a virtual property whose value is calculated based on the roles that have been granted to a user, either manually as in this example, or dynamically, through a script or condition. For more information about dynamically granted roles, see Section 9.4.3.2, "Granting Roles Dynamically" in the *Integrator's Guide*.

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/managed/user?_queryFilter=/givenName+eq+'Felicitas'&_fields=_id,userName,roles,effectiveRoles"
{
  "result" : [ {
    "_id" : "837085ae-766e-417c-9b7e-c36eee4352a3",
    "_rev" : "2",
    "userName" : "fdoe",
    "roles" : [ {
      "_ref": "managed/role/ad19979e-adbb-4d35-8320-6db50646b432",
      "_refProperties": {
        "_id": "4a42cd0b-d5d0-47e9-81e7-513aed74f6bc",
        "_rev": "1"
      }
    } ],
    "effectiveRoles" : [ {
      "_ref" : "managed/role/ad19979e-adbb-4d35-8320-6db50646b432"
    } ]
  } ]
},
...
}
```

Note that Felicitas's `roles` and `effectiveRoles` attributes both show the reference to the `Employee` role ID.

8.1.5. Removing a Managed User's Roles

Imagine that the `Employee` role was erroneously granted to Felicitas, and must be removed from her user entry.

To remove a granted role from a managed user, send a `DELETE` request to the user entry, specifying the ID of the `relationship` that must be removed. Note that this is not the ID of the role.

The request to remove the `Employee` role from Felicitas's entry is as follows:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request DELETE \
"http://localhost:8080/openidm/managed/user/837085ae-766e-417c-9b7e-c36eee4352a3/roles/4a42cd0b-d5d0-47e9-81e7-513aed74f6bc"
{
  "_ref": "managed/role/ad19979e-adbb-4d35-8320-6db50646b432",
  "_refProperties": {
    "_id": "4a42cd0b-d5d0-47e9-81e7-513aed74f6bc",
    "_rev": "1"
  }
}
```

If you query Felicitas's entry again, you will notice that her `roles` and `effectiveRoles` properties are now empty:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/managed/user?_queryFilter=/givenName+eq+'Felicitas'&_fields=_id,userName,roles,effectiveRoles"
{
  "result" : [ {
    "_id" : "837085ae-766e-417c-9b7e-c36eee4352a3",
    "_rev" : "2",
    "userName" : "fdoe",
    "roles": [],
    "effectiveRoles": []
  } ]
  ,
  ...
}
```

There are other methods to remove a granted role from a user over the REST interface. These are described in Section 9.4.6, "Deleting a User's Roles" in the *Integrator's Guide*.

Tip

You can also remove the granted role from fdoe's entry in the Admin UI as follows:

1. Select Manage > User and click on fdoe's entry.
2. On the Provisioning Roles tab, click the checkbox next to the Employee role and click Remove Selected Provisioning Roles.

8.1.6. Deleting a Managed Role

The final step in basic role management is to remove an existing role. In this section, we will remove the Contractor role we created previously.

Note that it is not possible to remove a role that is already granted to a user. To demonstrate this, we temporarily grant the Contractor role to Felicitas.

1. Assign the Contractor role to Felicitas, either by using the Admin UI, as described previously, or over the REST interface. If you use the REST interface, you will need to remember the system-generated identifier that was output when you queried the roles.

```
$ curl \
--header "Content-type: application/json" \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request PATCH \
--data '[
  {
    "operation" : "add",
    "field" : "/roles/-",
    "value" : {"_ref": "managed/role/b02d2531-5066-415e-bc90-31fe57e02322"}
  }
]' \
"http://localhost:8080/openidm/managed/user/837085ae-766e-417c-9b7e-c36eee4352a3"
{
  "_id": "837085ae-766e-417c-9b7e-c36eee4352a3",
  "_rev": "4",
  "mail": "fdoe@example.com",
  "sn": "Doe",
  "telephoneNumber": "555-1234",
  "userName": "fdoe",
  "givenName": "Felicitas",
  "description": "Felicitas Doe",
  "displayName": "fdoe",
  "accountStatus": "active",
  "effectiveRoles": [
    {
      "_ref": "managed/role/b02d2531-5066-415e-bc90-31fe57e02322"
    }
  ],
  "effectiveAssignments": []
}
```

2. Now, try to delete the Contractor role:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request DELETE \
"http://localhost:8080/openidm/managed/role/b02d2531-5066-415e-bc90-31fe57e02322"
{
  "code": 409,
  "reason": "Conflict",
  "message": "Cannot delete a role that is currently granted"
}
```

As you can see from the previous output, it is not possible to delete a role that is still granted to a user.

3. Remove the role from Felicitas's entry, either by using the Admin UI (as described previously), or by using a DELETE request, specifying the ID of the *relationship* that must be removed. Note that this is not the ID of the role. You will need to read Felicitas's `roles` property to obtain the relationship ID:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/managed/user?_queryFilter=/givenName+eq+'Felicitas'&_fields=_id,userName,roles,effectiveRoles"
{
  "result" : [ {
    "_id" : "837085ae-766e-417c-9b7e-c36eee4352a3",
    "_rev" : "3",
    "userName" : "fdoe",
    "roles" : [ {
      "_ref": "managed/role/b02d2531-5066-415e-bc90-31fe57e02322",
      "_refProperties": {
        "_id": "93703eff-a7ef-4cf3-80b1-f86fa3607978",
        "_rev": "1"
      }
    } ],
    "effectiveRoles" : [ {
      "_ref" : "managed/role/b02d2531-5066-415e-bc90-31fe57e02322"
    } ]
  } ]
},
...
}
```

In this example, the relationship ID is 93703eff-a7ef-4cf3-80b1-f86fa3607978.

The following request removes the Contractor role from Felicitas's entry:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request DELETE \
"http://localhost:8080/openidm/managed/user/837085ae-766e-417c-9b7e-c36eee4352a3/roles/93703eff-a7ef-4cf3-80b1-f86fa3607978"
{
  "_ref": "managed/role/b02d2531-5066-415e-bc90-31fe57e02322",
  "_refProperties": {
    "_id": "93703eff-a7ef-4cf3-80b1-f86fa3607978",
    "_rev": "1"
  }
}
```

4. Now that no users are granted the Contractor role, you can delete the role as follows:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request DELETE \
"http://localhost:8080/openidm/managed/role/b02d2531-5066-415e-bc90-31fe57e02322"
{
  "_id": "b02d2531-5066-415e-bc90-31fe57e02322",
  "_rev": "1",
  "name": "Contractor",
  "description": "Role granted to contract workers."
}
```

The DELETE operation returns the complete role entry.

Tip

Delete the Contractor role by using the Admin UI:

1. Select Manage > Role.
2. Select the Contractor role and click Delete Selected.

5. Verify that the Contractor role has been deleted by querying the list of managed role objects:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/managed/role?_queryFilter=true&_prettyPrint=true"
{
  "result": [
    {
      "_id": "ad19979e-adbb-4d35-8320-6db50646b432",
      "_rev": "1",
      "name": "Employee",
      "description": "Role granted to workers on the payroll."
    }
  ],
  ...
}
```

Note that only the Employee role remains.

Tip

List the remaining role objects in the Admin UI by clicking Manage > Role.

This concludes the basic role management operations. In the next section, you will see how to add assignments to roles, and how those assignments are used to provision users to external systems.

8.2. Provisioning Role Sample - Provisioning to an LDAP Server

The main purpose of OpenIDM roles is to provision a set of attributes, based on a managed user's role membership.

This sample builds on what you learnt in the previous sample, and you will create the same Employee and Contractor roles that were described in that sample. This sample also builds on Sample 2b (described in Section 3.2, "Sample 2b - LDAP Two Way"), and provisions users from the managed user repository to an OpenDJ directory.

The sample assumes a company, example.com. As an *Employee* of example.com, a user should be added to two groups in OpenDJ - the Employees group and the Chat Users group (presumably to access certain internal applications). As a *Contractor*, a user should be added only to the Contractors group in OpenDJ. A user's employee type must also be set correctly in OpenDJ, based on the role that is granted to the user.

8.2.1. External LDAP Configuration

Configure OpenDJ as for sample 2 (see Section 3.1.1, "LDAP Server Configuration"). The LDAP user must have write access to create users from OpenIDM on the LDAP server. When you set up the LDAP server, import the LDIF file for this sample ([openidm/samples/roles/provrole/data/Example.ldif](#)).

8.2.2. Before You Start

This section sets up the scenario by performing the following tasks:

1. Start OpenIDM with the configuration for the provisioning roles sample.
2. Create the Employee and Contractor roles that you created in the previous sample.
3. Reconcile the managed user repository with the user entries in the LDAP server.
1. Prepare OpenIDM as described in Section 1.3, "Preparing the Server", then start OpenIDM with the configuration for the Provisioning sample.

```
$ cd /path/to/openidm
```

```
$ startup.sh -p samples/roles/provrole
```

2. Create the Employee and Contractor roles, either by using the Admin UI (as described in the previous sample), or by running the following commands:

```
$ curl \
--header "Content-type: application/json" \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request POST \
--data '{
  "name" : "Employee",
  "description": "Role granted to workers on the payroll."
}' \
"http://localhost:8080/openidm/managed/role?_action=create"
{
  "_id" : "2902afd5-155a-49c0-9dd9-7e6bfcf1708a",
  "_rev" : "1",
  "name" : "Employee",
  "description" : "Role granted to workers on the payroll."
}
```

```
$ curl \
--header "Content-type: application/json" \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request POST \
--data '{
  "name" : "Contractor",
  "description": "Role granted to contract workers."
}' \
"http://localhost:8080/openidm/managed/role?_action=create"
{
  "_id": "e7f649ad-8013-4673-a52a-bdcac7483111",
  "_rev": "1",
  "name": "Contractor",
  "description": "Role granted to contract workers."
}
```

Note the IDs of these two roles because you will use them in the commands that follow.

3. Reconcile the repository.

The `sync.json` configuration file for this sample includes two mappings: `systemLdapAccounts_managedUser`, which synchronizes users from the source LDAP server with the target OpenIDM repository; and `managedUser_systemLdapAccounts`, which synchronizes changes from the OpenIDM repository with the LDAP server.

Run a reconciliation operation for the first mapping, either by using the Admin UI, or over the REST interface:

- To use the Admin UI, select Configure > Mapping, click on the first mapping (System/Ldap/Account --> Managed User) and click Reconcile Now.
- To use the REST interface, run the following command:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request POST \
"http://localhost:8080/openidm/recon?
action=recon&mapping=systemLdapAccounts_managedUser&waitForCompletion=true"
{
  "_id": "b5c535f8-5c1f-44dc-afa3-40d4f9984925-24",
  "state": "SUCCESS"
}
```

The sample is now ready to demonstrate provisioning roles.

8.2.3. Run the Sample

This section assumes that you have reconciled the managed user repository to populate it with the users from the LDAP server, and that you have created the Employee and Contractor roles.

This part of the sample demonstrates the following features of the OpenIDM roles implementation:

- Section 8.2.3.1, "Adding Assignments to a Role Definition"
- Section 8.2.3.2, "Granting a Role to a User and Observing that User's Role Assignments"
- Section 8.2.3.3, "Propagating Assignments to an External System"
- Section 8.2.3.4, "Removing a Role Grant From a User and Observing That User's Role Assignments"

8.2.3.1. Adding Assignments to a Role Definition

A role *assignment* is the logic that provisions a managed user to an external system, based on some criteria. The most common use case of a role assignment is the provisioning of specific attributes to an external system, based on the role or roles that the managed user has been granted. Assignments are sometimes called *entitlements*. For more information about assignments, see Section 9.4.8, "Working With Role Assignments" in the *Integrator's Guide*.

In this section, you will create assignments and add them to the two roles that you created previously. This section assumes the following scenario:

example.com's policy requires that every employee has the correct value for their `employeeType` in their corporate directory (OpenDJ).

1. Display the roles that you created in the previous section:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/managed/role?_queryFilter=true"
{
  "result" : [
    {
      "_id" : "2902afd5-155a-49c0-9dd9-7e6bfcf1708a",
      "_rev" : "1",
      "name" : "Employee",
      "description" : "Role granted to workers on the payroll."
    },
    {
      "_id" : "e7f649ad-8013-4673-a52a-bdcac7483111",
      "_rev" : "1",
      "name" : "Contractor",
      "description" : "Role granted to contract workers."
    }
  ]
},
...
}
```

Tip

Display the roles in the Admin UI by selecting Manage > Role.

2. Create a new managed assignment named Employee.

The assignment is specifically for the mapping from the managed user repository to the LDAP server. The assignment sets the value of the `employeeType` attribute on the LDAP server to `Employee`:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--header "Content-type: application/json" \
--request POST \
--data '{
  "name" : "Employee",
  "description": "Assignment for employees.",
  "mapping" : "managedUser_systemLdapAccounts",
  "attributes": [
    {
      "name": "employeeType",
      "value": "Employee",
      "assignmentOperation" : "mergeWithTarget",
      "unassignmentOperation" : "removeFromTarget"
    }
  ]
}' \
"http://localhost:8080/openidm/managed/assignment?_action=create"
{
  "_id": "f2830b80-6ab8-416b-b219-d3bf2efd0ed3",
  "_rev": "1",
```

```
"name": "Employee",
"description": "Assignment for employees.",
"mapping": "managedUser_systemLdapAccounts",
"attributes": [
  {
    "name": "employeeType",
    "value": "Employee",
    "assignmentOperation": "mergeWithTarget",
    "unassignmentOperation": "removeFromTarget"
  }
]
```

Tip

Create the assignment in the Admin UI:

1. Select Manage > Assignment, and click New Assignment.
2. Enter a name and description for the assignment, and select the mapping for which the assignment is applied (managedUser_systemLdapAccounts).
3. Click Add Assignment.
4. Select the Attributes tab, and click Add an Attribute.
5. Select employeeType, and enter the value Employee.
6. Click Save.

3. Add the assignment to the Employee role that you created previously.

Assignments are implemented as *relationship objects*. This means that you add an assignment to a role by *referencing* the assignment in the role's `assignments` field:

This command patches the Employee role to update its `assignments` field.

```
$ curl \
--header "Content-type: application/json" \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request PATCH \
--data '[
  {
    "operation" : "add",
    "field" : "/assignments/-",
    "value" : { "_ref": "managed/assignment/f2830b80-6ab8-416b-b219-d3bf2efd0ed3"}
  }
]' \
"http://localhost:8080/openidm/managed/role/2902afd5-155a-49c0-9dd9-7e6bfcf1708a"
{
  "_id": "2902afd5-155a-49c0-9dd9-7e6bfcf1708a",
  "_rev": "2",
  "name": "Employee",
  "description": "Role granted to workers on the payroll."
}
```

Tip

Add the assignment to the role in the Admin UI:

1. Select Manage > Role, and select the Employee role.
2. On the Managed Assignments tab, click Add Managed Assignments.
3. Select the Employee assignment and click Add.

8.2.3.2. Granting a Role to a User and Observing that User's Role Assignments

When a role is granted to a user (by updating the users `roles` property), any assignments that are referenced by the role are automatically referenced in the user's `assignments` property.

In this section, we will grant the Employee role we created previously to the user Barbara Jensen, who was created in the managed/user repository during the reconciliation from OpenDJ.

1. Before you can update Barbara Jensen's entry, determine the identifier of her entry by querying her username, `bjensen`, and requesting only her `_id` field:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/managed/user?_queryFilter=/userName+eq+'bjensen'&_fields=_id"

{
  "result" : [ {
    "_id" : "2c7daf46-d3ce-4bc5-9790-b44113bca8e7",
    "_rev" : "1"
  } ],
  ...
}
```

From the output, you can see that bjensen's `_id` is `2c7daf46-d3ce-4bc5-9790-b44113bca8e7`. (This unique ID will obviously be different in your command output.)

2. Update bjensen's entry by adding a reference to the ID of the Employee role as a value of her `roles` attribute:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--header "Content-type: application/json" \
--request PATCH \
--data '[
  {
    "operation" : "add",
    "field" : "/roles/-",
    "value" : { "_ref": "managed/role/2902afd5-155a-49c0-9dd9-7e6bfcf1708a" }
  }
]' \
"http://localhost:8080/openidm/managed/user/2c7daf46-d3ce-4bc5-9790-b44113bca8e7"

{
  "_id": "2c7daf46-d3ce-4bc5-9790-b44113bca8e7",
  "_rev": "4",
  "displayName": "Barbara Jensen",
  "description": "Created for OpenIDM",
  "givenName": "Barbara",
  "mail": "bjensen@example.com",
  "telephoneNumber": "1-360-229-7105",
  "sn": "Jensen",
  "userName": "bjensen",
  "accountStatus": "active",
  "effectiveRoles": [
    {
      "_ref": "managed/role/2902afd5-155a-49c0-9dd9-7e6bfcf1708a"
    }
  ],
  "effectiveAssignments": [
    {
      "name": "Employee",
      "description": "Assignment for employees.",
      "mapping": "managedUser_systemLdapAccounts",
      "attributes": [
        {
          "name": "employeeType",
```

```

    "value": "Employee",
    "assignmentOperation": "mergeWithTarget",
    "unassignmentOperation": "removeFromTarget"
  }
],
"_id": "f2830b80-6ab8-416b-b219-d3bf2efd0ed3",
"_rev": "1"
}
]
}

```

Tip

Assign the role to bjensen by using the Admin UI:

1. Select Manage > User, and click on bjensen's entry.
2. On the Provisioning Roles tab, click Add Provisioning Roles.
3. Select the Employee role and click Add.

3. Take a closer look at bjensen's entry, specifically at her roles, effective roles and effective assignments:

```

$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/managed/user?_queryFilter=/userName+eq+'bjensen'&_fields=_id,userName,roles,effectiveRoles,effectiveAssignments"
{
  "result": [
    {
      "_id": "2c7daf46-d3ce-4bc5-9790-b44113bca8e7",
      "_rev": "4",
      "userName": "bjensen",
      "roles": [
        {
          "_ref": "managed/role/2902afd5-155a-49c0-9dd9-7e6bfcf1708a",
          "_refProperties": {
            "_id": "b1c29213-e726-466a-9051-e9bb4e593331",
            "temporalConstraints": [],
            "_grantType": "",
            "_rev": "2"
          }
        }
      ]
    },
    {
      "effectiveRoles": [
        {
          "_ref": "managed/role/2902afd5-155a-49c0-9dd9-7e6bfcf1708a"
        }
      ],
      "effectiveAssignments": [
        {
          "name": "Employee",

```



```
"description": "Assignment for employees.",
"mapping": "managedUser_systemLdapAccounts",
"attributes": [
  {
    "name": "employeeType",
    "value": "Employee",
    "assignmentOperation": "mergeWithTarget",
    "unassignmentOperation": "removeFromTarget"
  }
],
"_id": "f2830b80-6ab8-416b-b219-d3bf2efd0ed3",
"_rev": "1"
}
]
}
],
...
}
```

Note that bjensen now has the calculated property `effectiveAssignments`, which includes the set of assignments that pertains to any user with the Employee role. Currently the assignment lists the `employeeType` attribute.

In the next section, you will see how the assignment is used to set the value of the `employeeType` attribute in the LDAP server.

8.2.3.3. Propagating Assignments to an External System

This section provides a number of steps that show how effective assignments are propagated out to the external systems associated with their mappings.

1. Verify that bjensen's `employeeType` has been set correctly in the OpenDJ.

Because implicit synchronization is enabled by default in OpenIDM, any changes made to a managed user object are pushed out to all the external systems for which mappings are configured.

So, because bjensen has an effective assignment that sets an attribute in her LDAP entry, you should immediately see the resulting change in her LDAP entry.

To verify that her entry has changed, run an `ldapsearch` on her entry and check the value of her `employeeType` attribute.

Note

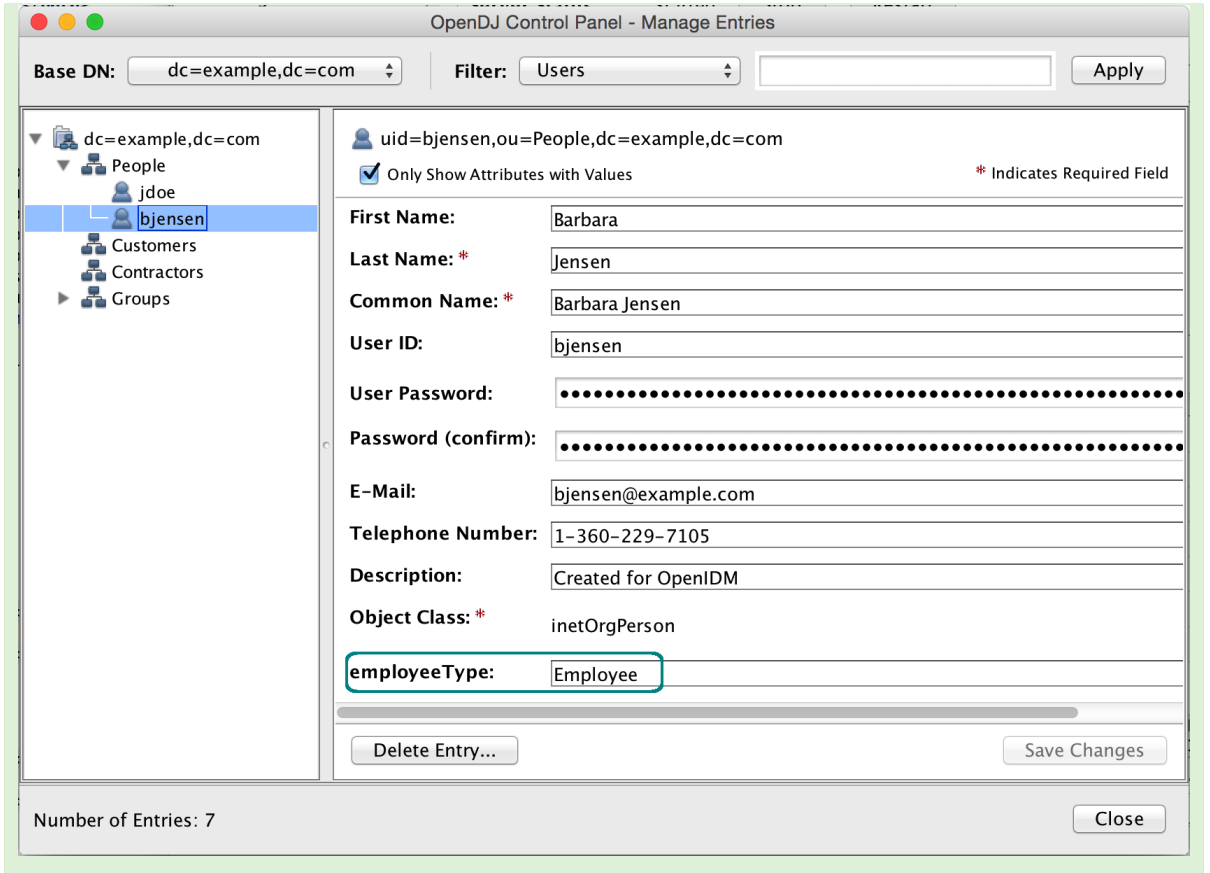
This command assumes that you are using the `ldapsearch` provided with OpenDJ (`opendj/bin/ldapsearch`).

```
$ ldapsearch \  
--port 1389 \  
--hostname localhost \  
--baseDN "dc=example,dc=com" \  
--bindDN "cn=Directory Manager" \  
--bindPassword password \  
--searchScope sub \  
"(uid=bjensen)" dn uid employeeType  
dn: uid=bjensen,ou=People,dc=example,dc=com  
uid: bjensen  
employeeType: Employee
```

Note that bjensen's `employeeType` attribute is correctly set to `Employee`.

Tip

You can also check bjensen's LDAP entry by using the OpenDJ Control Panel (`opendj/bin/control-panel`):



- To observe how a managed user's roles can be used to provision group membership in an external directory, we add the groups that an Employee and a Contractor should have in the corporate directory (OpenDJ) as assignment attributes of the respective roles.

First, look at the current **assignments** of the Employee role again:

```

$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/managed/role/2902afd5-155a-49c0-9dd9-7e6bfcf1708a?_fields=assignments
,name"
{
  "_id": "2902afd5-155a-49c0-9dd9-7e6bfcf1708a",
  "_rev": "2",
  "assignments": [
    {
      "_ref": "managed/assignment/f2830b80-6ab8-416b-b219-d3bf2efd0ed3",
      "_refProperties": {
        "_id": "c0005ecb-9dda-4db1-8660-a723b8237f16",
        "temporalConstraints": [],
        "_grantType": "",
        "_rev": "1"
      }
    }
  ],
  "name": "Employee"
}

```

To update the `groups` attribute in bjensen's LDAP entry, you do not need to create a *new* assignment. You simply need to add the attribute for LDAP groups to the assignment with ID `f2830b80-6ab8-416b-b219-d3bf2efd0ed3`:

```

$ curl \
--header "Content-type: application/json" \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request PATCH \
--data '[
  {
    "operation" : "add",
    "field" : "/attributes/-",
    "value" : {
      "name": "ldapGroups",
      "value": [
        "cn=Employees,ou=Groups,dc=example,dc=com",
        "cn=Chat Users,ou=Groups,dc=example,dc=com"
      ],
      "assignmentOperation" : "mergeWithTarget",
      "unassignmentOperation" : "removeFromTarget"
    }
  }
]' \
"http://localhost:8080/openidm/managed/assignment/f2830b80-6ab8-416b-b219-d3bf2efd0ed3"

{
  "_id": "f2830b80-6ab8-416b-b219-d3bf2efd0ed3",
  "_rev": "2",
  "name": "Employee",
  "description": "Assignment for employees.",
  "mapping": "managedUser_systemLdapAccounts",
  "attributes": [

```

```
{
  "name": "employeeType",
  "value": "Employee",
  "assignmentOperation": "mergeWithTarget",
  "unassignmentOperation": "removeFromTarget"
},
{
  "name": "ldapGroups",
  "value": [
    "cn=Employees,ou=Groups,dc=example,dc=com",
    "cn=Chat Users,ou=Groups,dc=example,dc=com"
  ],
  "assignmentOperation": "mergeWithTarget",
  "unassignmentOperation": "removeFromTarget"
}
]
```

So, the Employee assignment now sets two attributes on the LDAP system - the `employeeType` attribute, and the `ldapGroups` attribute.

Tip

To add more attributes to the Employee assignment in the Admin UI:

1. Select Manage > Assignment, and click on the Employee assignment.
2. On the Attributes tab, select Add an attribute and select the `ldapGroups` attribute.
3. Enter the values

```
cn=Employees,ou=Groups,dc=example,dc=com
```

and

```
cn=Chat Users,ou=Groups,dc=example,dc=com
```

and click Save.

3. With the implicit synchronization between the managed user repository and OpenDJ, `bjensen` should now be a member of the `cn=Employees` and `cn=Chat Users` groups in LDAP.

You can verify this with the following `ldapsearch` command. This command returns `bjensen`'s group membership, in her `isMemberOf` attribute.

```
$ ldapsearch \
--port 1389 \
--hostname localhost \
--baseDN "dc=example,dc=com" \
--bindDN "cn=Directory Manager" \
--bindPassword password \
--searchScope sub \
"(uid=bjensen)" dn uid employeeType isMemberOf

dn: uid=bjensen,ou=People,dc=example,dc=com
uid: bjensen
employeeType: Employee
isMemberOf: cn=openidm2,ou=Groups,dc=example,dc=com
isMemberOf: cn=Chat Users,ou=Groups,dc=example,dc=com
isMemberOf: cn=Employees,ou=Groups,dc=example,dc=com
```

You can also check bjensen's group membership by using the OpenDJ Control Panel (as shown previously), or by querying her object in the LDAP system, over the REST interface:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/system/ldap/account?_queryFilter=/uid+sw+'bjensen'&_fields=dn,uid,employeeType,ldapGroups"
{
  "result" : [ {
    "_id" : "uid=bjensen,ou=People,dc=example,dc=com",
    "dn" : "uid=bjensen,ou=People,dc=example,dc=com",
    "uid" : "bjensen",
    "employeeType" : "Employee",
    "ldapGroups" : [
      "cn=openidm2,ou=Groups,dc=example,dc=com",
      "cn=Employees,ou=Groups,dc=example,dc=com",
      "cn=Chat Users,ou=Groups,dc=example,dc=com"
    ]
  } ],
  ...
}
```

In the original LDIF file, bjensen was already a member of the openidm2 group. You can ignore this group for the purposes of this sample.

Tip

Use the Admin UI to see bjensen's LDAP groups as follows:

1. Select Manage > User, and select bjensen's entry.
2. On the Linked Systems tab, scroll down to the ldapGroups item.

4. Now, create a new assignment that will apply to Contract employees, and add that assignment to the Contractor role.

Create the Contractor assignment with the following command. This assignment sets the value of the `employeeType` attribute to `Contractor`, and updates the user's `ldapGroups` attribute to include the `cn=Contractors` group:

```
$ curl \
--header "Content-type: application/json" \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request POST \
--data '{
  "name": "Contractor",
  "description": "Contractor assignment for contract workers.",
  "mapping": "managedUser_systemLdapAccounts",
  "attributes": [
    {
      "name": "ldapGroups",
      "value": [
        "cn=Contractors,ou=Groups,dc=example,dc=com"
      ],
      "assignmentOperation": "mergeWithTarget",
      "unassignmentOperation": "removeFromTarget"
    },
    {
      "name": "employeeType",
      "value": "Contractor",
      "assignmentOperation": "mergeWithTarget",
      "unassignmentOperation": "removeFromTarget"
    }
  ]
}' \
"http://localhost:8080/openidm/managed/assignment?_action=create"
{
  "_id": "7536e234-1268-482d-8459-24c8ef832def",
  "_rev": "1",
  "name": "Contractor",
  "description": "Contractor assignment for contract workers.",
  "mapping": "managedUser_systemLdapAccounts",
  "attributes": [
    {
      "name": "ldapGroups",
      "value": [
        "cn=Contractors,ou=Groups,dc=example,dc=com"
      ],
      "assignmentOperation": "mergeWithTarget",
      "unassignmentOperation": "removeFromTarget"
    },
    {
      "name": "employeeType",
      "value": "Contractor",
      "assignmentOperation": "mergeWithTarget",
      "unassignmentOperation": "removeFromTarget"
    }
  ]
}
```

Note the ID of the Contractor assignment (7536e234-1268-482d-8459-24c8ef832def in this example).

Tip

Create the assignment by using the Admin UI, as described in Section 8.2.3.1, "Adding Assignments to a Role Definition").

5. Now, add the Contractor assignment to the Contractor role (ID e7f649ad-8013-4673-a52a-bdcac7483111 in this example):

```
$ curl \
--header "Content-type: application/json" \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request PATCH \
--data '[
  {
    "operation" : "add",
    "field" : "/assignments/-",
    "value" : {
      "_ref" : "managed/assignment/7536e234-1268-482d-8459-24c8ef832def"
    }
  }
]' \
"http://localhost:8080/openidm/managed/role/e7f649ad-8013-4673-a52a-bdcac7483111"
{
  "_id": "e7f649ad-8013-4673-a52a-bdcac7483111",
  "_rev": "2",
  "name": "Contractor",
  "description": "Role granted to contract workers."
}
```

Tip

Add the Contractor assignment to the Contractor role in the Admin UI, as follows:

1. Select Manage > Role, and select the Contractor role.
2. On the Managed Assignments tab, click Add Managed Assignment.
3. Select the Contractor assignment from the dropdown list, and click Add.

6. Next, we need to grant the Contractor role to user jdoe. Before we can patch jdoe's entry, we need to know his system-generated ID. To obtain the ID, query jdoe's entry as follows:


```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/managed/user?_queryFilter=/userName+eq+'jdoe'&_fields=_id"

{
  "result": [
    {
      "_id": "92680be0-82f9-4297-9e00-c35c7cf700d2",
      "_rev": "2"
    }
  ],
  ...
}
```

From the output, you can see that jdoe's `_id` is `92680be0-82f9-4297-9e00-c35c7cf700d2`. (This unique ID will obviously be different in your command output.)

7. Update jdoe's entry by adding a reference to the ID of the Contractor role (`e7f649ad-8013-4673-a52a-bdcac7483111`) as a value of his `roles` attribute:

```
$ curl \
--header "Content-type: application/json" \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request PATCH \
--data '[
  {
    "operation": "add",
    "field": "/roles/-",
    "value": {
      "_ref": "managed/role/e7f649ad-8013-4673-a52a-bdcac7483111"
    }
  }
]' \
"http://localhost:8080/openidm/managed/user/92680be0-82f9-4297-9e00-c35c7cf700d2"

{
  "_id": "92680be0-82f9-4297-9e00-c35c7cf700d2",
  "_rev": "4",
  "displayName": "John Doe",
  "description": "Created for OpenIDM",
  "givenName": "John",
  "mail": "jdoe@example.com",
  "telephoneNumber": "1-415-599-1100",
  "sn": "Doe",
  "userName": "jdoe",
  "accountStatus": "active",
  "effectiveRoles": [
    {
      "_ref": "managed/role/e7f649ad-8013-4673-a52a-bdcac7483111"
    }
  ],
  "effectiveAssignments": [
    {
      "name": "Contractor",
      "description": "Contractor assignment for contract workers.",
    }
  ]
}
```

```
"mapping": "managedUser_systemLdapAccounts",
"attributes": [
  {
    "name": "ldapGroups",
    "value": [
      "cn=Contractors,ou=Groups,dc=example,dc=com"
    ],
    "assignmentOperation": "mergeWithTarget",
    "unassignmentOperation": "removeFromTarget"
  },
  {
    "name": "employeeType",
    "value": "Contractor",
    "assignmentOperation": "mergeWithTarget",
    "unassignmentOperation": "removeFromTarget"
  }
],
"_id": "7536e234-1268-482d-8459-24c8ef832def",
"_rev": "1"
}
]
```

Tip

Grant the Contractor role to jdoe by using the Admin UI, as follows:

1. Select Manage > User, and click on jdoe's entry.
2. On the Provisioning Roles tab, click Add Provisioning Roles.
3. Select the Contractor role and click Add.
4. Click Save.

8. Check jdoe's entry on the LDAP system.

With the implicit synchronization between the managed user repository and OpenDJ, jdoe should now be a member of the `cn=Contractors` group in LDAP. In addition, his `employeeType` should have been set to `Contractor`.

You can verify this with the following REST query. This command returns jdoe's group membership, in his `isMemberOf` attribute, and his `employeeType`:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/system/ldap/account?_queryFilter=/uid+sw+'jdoe'&_fields=dn,uid,employeeType,ldapGroups"
{
  "result": [
    {
      "_id": "uid=jdoe,ou=People,dc=example,dc=com",
      "givenName": "John",
      "ldapGroups": [
        "cn=openidm,ou=Groups,dc=example,dc=com",
        "cn=Contractors,ou=Groups,dc=example,dc=com"
      ],
      "mail": "jdoe@example.com",
      "employeeType": "Contractor",
      "uid": "jdoe",
      "telephoneNumber": "1-415-599-1100",
      "sn": "Doe",
      "disabled": null,
      "cn": "John Doe",
      "description": "Created for OpenIDM",
      "dn": "uid=jdoe,ou=People,dc=example,dc=com"
    }
  ],
  ...
}
```

Tip

Use the Admin UI to see jdoe's LDAP groups as follows:

1. Select Manage > User, and select jdoe's entry.
2. On the Linked Systems tab, scroll down to the ldapGroups item.

8.2.3.4. Removing a Role Grant From a User and Observing That User's Role Assignments

In this section, you will remove the Contractor role from jdoe's managed user entry and observe the subsequent change to jdoe's managed assignments, and to the corresponding attributes in OpenDJ.

1. Before you change jdoe's roles, view his entry again to examine his current roles.

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/managed/user?_queryFilter=/userName+eq+'jdoe'&_fields=_id,roles"
{
  "result": [
    {
      "_id": "92680be0-82f9-4297-9e00-c35c7cf700d2",
      "_rev": "4",
      "roles": [
        {
          "_ref": "managed/role/e7f649ad-8013-4673-a52a-bdcac7483111",
          "_refProperties": {
            "_id": "093fc34b-0694-478e-952e-98d0a828b1ac",
            "_rev": "2"
          }
        }
      ]
    }
  ],
  ...
}
```

Note that jdoe's ID is 92680be0-82f9-4297-9e00-c35c7cf700d2 and the ID of the *relationship* that expresses the role grant is 093fc34b-0694-478e-952e-98d0a828b1ac. You will need these IDs in the next step.

Tip

View jdoe's current roles in the Admin UI:

1. Select Manage > User, and select jdoe's entry.
2. The Provisioning Roles tab lists the current roles.

2. Remove the Contractor role from jdoe's entry by sending a DELETE request to the user entry, specifying the relationship ID:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request DELETE \
"http://localhost:8080/openidm/managed/user/92680be0-82f9-4297-9e00-c35c7cf700d2/roles/093fc34b-0694-478e-952e-98d0a828b1ac"
{
  "_ref": "managed/role/e7f649ad-8013-4673-a52a-bdcac7483111",
  "_refProperties": {
    "_id": "093fc34b-0694-478e-952e-98d0a828b1ac",
    "_rev": "2"
  }
}
```

Tip

Use the Admin UI to remove the Contractor role from jdoe's entry as follows:

1. Select Manage > User, and select jdoe's entry.
2. On the Provisioning Roles tab, check the box next to the Contractor role and click Remove Selected Provisioning Roles.

3. Verify jdoe's `employeeType` and `ldapGroups`.

The removal of the Contractor role causes a synchronization operation to be run on jdoe's entry. His `employeeType` and `ldapGroups` attributes in OpenDJ should be reset to what they were before he was granted the Contractor role.

You can check jdoe's attributes by querying his object in the LDAP directory, over the REST interface:

```
$ curl \
  --header "X-OpenIDM-Username: openidm-admin" \
  --header "X-OpenIDM-Password: openidm-admin" \
  --request GET \
  "http://localhost:8080/openidm/system/ldap/account?_queryFilter=/uid+sw+'jdoe'&_fields=dn,uid,employeeType,ldapGroups"
{
  "result" : [ {
    "sn" : "Doe",
    "telephoneNumber" : "1-415-599-1100",
    "employeeType" : null,
    "dn" : "uid=jdoe,ou=People,dc=example,dc=com",
    "cn" : "John Doe",
    "uid" : "jdoe",
    "ldapGroups" : [ "cn=openidm,ou=Groups,dc=example,dc=com" ],
    "givenName" : "John",
    "mail" : "jdoe@example.com",
    "description" : "Created for OpenIDM",
    "_id" : "uid=jdoe,ou=People,dc=example,dc=com"
  } ],
  ...
}
```

Tip

Use the Admin UI to see jdoe's LDAP groups as follows:

1. Select Manage > User, and select jdoe's entry.

2. On the Linked Systems tab, scroll down to the ldapGroups item.

This concludes the provisioning with roles sample. For more information about roles, assignments, and how to manipulate them, see Section 9.4, "Working With Managed Roles" in the *Integrator's Guide*.

8.3. Temporal Constraints Sample - Applying Time-Based Constraints to Roles

To restrict the period during which a role is effective, you can set a *temporal constraint* on the role itself, or on the role grant. A temporal constraint that is set on a role definition applies to all grants of that role. A temporal constraint that is set on a role grant enables you to specify the period that the role is valid per user. For more information about temporal constraints, see Section 9.4.4, "Using Temporal Constraints to Restrict Effective Roles" in the *Integrator's Guide*.

This sample is new in OpenIDM 4.5.0 and builds on what you learned in the previous two roles samples. The sample demonstrates how to add a temporal constraint to the Contractor role to restrict the period that it is effective. The previous sample ended with jdoe's contract ending because he finished up the work that was required for the company. The company now has additional work that requires jdoe's expertise and has rehired him as a contractor. This time however, the company requires that jdoe's contractor role be terminated as soon as his contract expires.

8.3.1. Before You Start

This sample assumes that you have already run through the previous two roles samples, and have installed and configured OpenDJ as for Sample 2 (see Section 3.1.1, "LDAP Server Configuration"). The LDAP user must have write access to create users from OpenIDM on the LDAP server. When you set up the LDAP server, import the LDIF file for this sample ([openidm/samples/roles/temporalConstraints/data/Example.ldif](#)).

1. Prepare OpenIDM as described in Section 1.3, "Preparing the Server", then start OpenIDM with the configuration for the temporal constraints sample.

```
$ cd /path/to/openidm
$ startup.sh -p samples/roles/temporalConstraints
```

2. Create the Employee and Contractor roles, and their assignments, either by using the Admin UI, or by using the REST interface, as described in the previous sample

Note the IDs of the roles and assignments because you will use them in the commands that follow.

3. Reconcile the repository, as described in the previous sample.

The sample is now ready to demonstrate temporal constraints on roles.

8.3.2. Run the Sample

This section assumes that you have reconciled the managed user repository to populate it with the users from the LDAP server, and that you have created the Employee and Contractor roles and their assignments.

This part of the sample demonstrates the following features of the OpenIDM roles implementation:

- Section 8.3.2.1, "Adding a Temporal Constraint to a Role Definition"
- Section 8.3.2.2, "Granting a Role With a Temporal Constraint to a User"
- Section 8.3.2.3, "Observing the Effects of a Temporal Role"
- Section 8.3.2.4, "Removing the Role Grant From the User"

8.3.2.1. Adding a Temporal Constraint to a Role Definition

In this section, you will update the definition of the Contractor role to add a temporal constraint.

1. If you are running the sample directly over the REST interface, obtain the ID of the Contractor role:

You can skip this step if you are using the Admin UI.

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/managed/role?_queryFilter=true&_prettyPrint=true"
{
  "result" : [ {
    "_id" : "1a4c9047-1ce3-4f39-8901-e4f60176330e",
    "_rev" : "1",
    "name" : "Employee",
    "description" : "Role granted to workers on the payroll."
  }, {
    "_id" : "06128fc1-b89b-4fe8-b9b3-4de30fd17b9e",
    "_rev" : "1",
    "name" : "Contractor",
    "description" : "Role granted to contract workers."
  } ],
  ...
}
```

In this example, the ID of the Contractor role is 06128fc1-b89b-4fe8-b9b3-4de30fd17b9e.

2. Patch the contractor role with a temporal constraint by adding a start and end date to the role definition.

For the start date, use the current time plus 5 minutes. This will demonstrate that you can add a future temporal constraint on a user before their contract even starts. For the end date, use 10 minutes from the current time.

In this example, the current time is 15:30:00. Adding 5 minutes results in a start time of 15:35:00. Adding 10 minutes results in an end time of 15:40:00. Adjust these values for the current time that you run the sample.

To add the temporal constraint over REST:

```
$ curl \
--header "Content-type: application/json" \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request PATCH \
--data '[
  {
    "operation" : "add",
    "field" : "/temporalConstraints",
    "value" : [{"duration": "2016-05-05T15:35:00.000Z/2016-05-05T15:40:00.000Z"}]
  }
]' \
"http://localhost:8080/openidm/managed/role/06128fc1-b89b-4fe8-b9b3-4de30fd17b9e"
{
  "_id": "06128fc1-b89b-4fe8-b9b3-4de30fd17b9e",
  "_rev": "2",
  "name": "Contractor",
  "description": "Role granted to contract workers.",
  "temporalConstraints": [
    {
      "duration": "2016-05-05T15:35:00.000Z/2016-05-05T15:40:00.000Z"
    }
  ]
}
```

To add the temporal constraint by using the Admin UI:

1. Select Manage > Role and click the Contractor role.
2. Select Temporal Constraint, then select a time zone, start date, and end date.
3. Click Save.

8.3.2.2. Granting a Role With a Temporal Constraint to a User

In this section, you will grant the Contractor role to jdoe.

1. If you are running this sample over REST, determine the identifier of jdoe's entry querying the existing managed users.

If you are using the Admin UI you can skip this step.


```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/managed/user?_queryFilter=true"
{
  "result" : [ {
    ...
  }, {
    "_id" : "0c470c71-bb4e-4cf1-bc9d-77d7b35b180b",
    "_rev" : "2",
    "displayName" : "John Doe",
    "description" : "Created for OpenIDM",
    "givenName" : "John",
    "mail" : "jdoe@example.com",
    "telephoneNumber" : "1-415-599-1100",
    "sn" : "Doe",
    "userName" : "jdoe",
    "accountStatus" : "active",
    "effectiveRoles" : [ ],
    "effectiveAssignments" : [ ]
  } ],
  ...
}
```

Note that jdoe's ID is 0c470c71-bb4e-4cf1-bc9d-77d7b35b180b.

2. Update jdoe's entry to grant him the Contractor role:

```
$ curl \
--header "Content-type: application/json" \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request PATCH \
--data '[
  {
    "operation" : "add",
    "field" : "/roles/-",
    "value" : {
      "_ref" : "managed/role/06128fc1-b89b-4fe8-b9b3-4de30fd17b9e"
    }
  }
]' \
"http://localhost:8080/openidm/managed/user/0c470c71-bb4e-4cf1-bc9d-77d7b35b180b"
{
  "_id": "0c470c71-bb4e-4cf1-bc9d-77d7b35b180b",
  "_rev": "4",
  "displayName": "John Doe",
  "description": "Created for OpenIDM",
  "givenName": "John",
  "mail": "jdoe@example.com",
  "telephoneNumber": "1-415-599-1100",
  "sn": "Doe",
  "userName": "jdoe",
  "accountStatus": "active",
  "effectiveRoles": [
  ],

```

```
"effectiveAssignments": [  
  ]  
}
```

Note that, at this stage, jdoe does not have any effective roles or assignments because his contract has not yet started. Check jdoe's entry again between the start and end date of the temporal constraint to observe the assignment.

To grant the Contractor role to jdoe by using the Admin UI:

1. Select Manage > User and click jdoe's entry.
2. On the Provisioning Roles tab, click Add Provisioning Roles.
3. Select the Contractor role and click Add.

8.3.2.3. Observing the Effects of a Temporal Role

During the contract period, query jdoe again to see when the Contractor role and assignments are present in his list of effective roles and effective assignments.

1. When the start time that you set has been reached, query jdoe's entry as follows:

```
$ curl \br/>  --header "X-OpenIDM-Username: openidm-admin" \  
  --header "X-OpenIDM-Password: openidm-admin" \  
  --request GET \  
  "http://localhost:8080/openidm/managed/user?_queryFilter=/userName+sw+"jdoe"&_fields=roles  
,effectiveRoles,effectiveAssignments"  
{  
  "result": [  
    {  
      "_id": "0c470c71-bb4e-4cf1-bc9d-77d7b35b180b",  
      "_rev": "6",  
      "roles": [  
        {  
          "_ref": "managed/role/06128fc1-b89b-4fe8-b9b3-4de30fd17b9e",  
          "_refProperties": {  
            "_id": "6774696d-999e-4483-87a9-02f501752b29",  
            "_rev": "4"  
          }  
        }  
      ],  
      "effectiveRoles": [  
        {  
          "_ref": "managed/role/06128fc1-b89b-4fe8-b9b3-4de30fd17b9e"  
        }  
      ],  
      "effectiveAssignments": [  
        {  
          "name": "Contractor",  
          "description": "Contractor assignment for contract workers.",  
          "mapping": "managedUser_systemLdapAccounts",  
          "attributes": [  

```

```

    {
      "name": "ldapGroups",
      "value": [
        "cn=Contractors,ou=Groups,dc=example,dc=com"
      ],
      "assignmentOperation": "mergeWithTarget",
      "unassignmentOperation": "removeFromTarget"
    },
    {
      "name": "employeeType",
      "value": "Contractor",
      "assignmentOperation": "mergeWithTarget",
      "unassignmentOperation": "removeFromTarget"
    }
  ],
  "_id": "3e61a1db-202d-4761-aeb7-b617d3376a79",
  "_rev": "1"
}
]
}
...
}

```

There is not really a comparable way to perform this step in the UI.

- Between the start and end times, look at `jdoue`'s entry in the LDAP server. At this point, the Contractor group should be effective for `jdoue`'s entry, so his LDAP groups should include the Contractors group:

```

$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/system/ldap/account?_queryFilter=/uid+sw+"jdoue"&_fields=dn,uid,employeeType,ldapGroups"
{
  "result": [
    {
      "_id": "uid=jdoue,ou=People,dc=example,dc=com",
      "dn": "uid=jdoue,ou=People,dc=example,dc=com",
      "uid": "jdoue",
      "employeeType": "Contractor",
      "ldapGroups": [
        "cn=openidm,ou=Groups,dc=example,dc=com",
        "cn=Contractors,ou=Groups,dc=example,dc=com"
      ]
    }
  ]
}

```

To look at `jdoue`'s LDAP groups in the Admin UI:

- Select Manage > User, and select `jdoue`'s entry.
- On the Linked Systems tab, scroll down to the `ldapGroups` item.

3. Jdoe's contract expires when the current date is greater than the end date of the temporal constraint.

Query jdoe's entry at that point to see that he no longer has the Contractor roles and assignments in his effective roles and assignments lists (although the role is still there in his roles list, it is no longer effective).

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/managed/user?_queryFilter=/userName+sw+"jdoe"&_fields=roles,
effectiveRoles,effectiveAssignments"
{
  "result": [
    {
      "_id": "0c470c71-bb4e-4cf1-bc9d-77d7b35b180b",
      "_rev": "8",
      "roles": [
        {
          "_ref": "managed/role/06128fc1-b89b-4fe8-b9b3-4de30fd17b9e",
          "_refProperties": {
            "_id": "6774696d-999e-4483-87a9-02f501752b29",
            "_rev": "6"
          }
        }
      ]
    },
    "effectiveRoles": [
  ],
    "effectiveAssignments": [
  ]
  }
],
  ...
}
```

There is not really comparable way to view this information in the Admin UI. Take note of the ID of the *relationship* that expresses the role grant (6774696d-999e-4483-87a9-02f501752b29 in this example). You will need this ID in the next step.

4. Now verify that jdoe is no longer part of the Contractors group in OpenDJ:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/system/ldap/account?_queryFilter=/uid+sw+"jdoe"&_fields=dn,uid,employeeType,ldapGroups"
{
  "result": [
    {
      "_id": "uid=jdoe,ou=People,dc=example,dc=com",
      "dn": "uid=jdoe,ou=People,dc=example,dc=com",
      "uid": "jdoe",
      "employeeType": null,
      "ldapGroups": [
        "cn=openidm,ou=Groups,dc=example,dc=com"
      ]
    }
  ],
  ...
}
```

In the Admin UI:

1. Select Manage > User, and select jdoe's entry.
2. On the Linked Systems tab, scroll down to the ldapGroups item.

8.3.2.4. Removing the Role Grant From the User

To finish up this sample, remove the Contractor role from jdoe by sending a DELETE request to his managed user entry, specifying the relationship ID (6774696d-999e-4483-87a9-02f501752b29 in this example):

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request DELETE \
"http://localhost:8080/openidm/managed/user/0c470c71-bb4e-4cf1-bc9d-77d7b35b180b/roles/6774696d-999e-4483-87a9-02f501752b29"
{
  "_ref": "managed/role/06128fc1-b89b-4fe8-b9b3-4de30fd17b9e",
  "_refProperties": {
    "_id": "6774696d-999e-4483-87a9-02f501752b29",
    "_rev": "6"
  }
}
```

Alternatively, use the Admin UI to remove the Contractor role from jdoe's entry as follows:

1. Select Manage > User, and select jdoe's entry.
2. On the Provisioning Roles tab, check the box next to the Contractor role and click Remove Selected Provisioning Roles.

Chapter 9

The Multi-Account Linking Sample

The sample provided in the `samples/multiaccountlinking` directory illustrates how OpenIDM addresses links from multiple accounts to a single identity.

This sample is based on a common use case in the insurance industry, where a company (Example.com) employs agents to sell policies to their insured customers. Most of their agents are also insured, which means that they have two distinct *roles* within the company - customers, and agents. With minor changes, this sample works for other use cases. For example, you might have a hospital that employs doctors who treat patients. Some of the doctors are also patients of that hospital.

9.1. Before You Start: Link Qualifiers, Agents, and Insured Customers

You define mappings between source and target accounts in the your project's `sync.json` file. As part of a mapping, you can create a link between a single source account and multiple target accounts using a *link qualifier*, that enables one-to-many relationships in mappings and policies. For more information about mappings and link qualifiers, see Section 14.3.2, "Mapping Source Objects to Target Objects" in the *Integrator's Guide* and Section 14.3.2.5, "Mapping a Single Source Object to Multiple Target Objects" in the *Integrator's Guide*.

This sample uses two link qualifiers:

- **Insured** represents the accounts associated with Example.com's Insured customers, created under the LDAP container `ou=Customers,dc=example,dc=com`.
- **Agent** represents agent accounts, considered independent contractors, and created under the LDAP container `ou=Contractors,dc=example,dc=com`.

Assume that agents and insured customers connect via two different portals, and that each group has access to different features, depending on the portal.

Agents might have two separate accounts; one each for professional and personal use. Although the accounts are different, the identity information for each agent should be the same for both accounts.

This sample therefore uses link qualifiers to distinguish the two *categories* of users. The link qualifiers are named **insured** and **agent**, and are defined as part of the `managedUser_systemLdapAccounts` mapping in the `sync.json` file:

```
{
  "name" : "managedUser_systemLdapAccounts",
  "source" : "managed/user",
  "target" : "system/ldap/account",
  "linkQualifiers" : [
    "insured",
    "agent"
  ],
  .....
}
```

You can check this configuration in the Admin UI. Click **Configure > Mappings > managedUser_systemLdapAccounts > Properties > Link Qualifiers**. You should see **insured** and **agent** in the list of configured link qualifiers.

In addition, the sample uses a transformation script that determines the LDAP Distinguished Name (**dn**) from the user category. The following excerpt of the `sync.json` file shows that script:

```
{
  "target" : "dn",
  "transform" : {
    "type" : "text/javascript",
    "globals" : { },
    "source" :
      "if (linkQualifier === 'agent') {
        'uid=' + source.userName + ',ou=Contractors,dc=example,dc=com';
      } else if (linkQualifier === 'insured') {
        'uid=' + source.userName + ',ou=Customers,dc=example,dc=com';
      }"
  },
}
```

Finally, the following `validSource` script assesses the effective roles of a managed user to determine if that user has an **Agent** or **Insured** role. The script then assigns a link qualifier based on the assessed role.

```
"validSource" : {
  "type" : "text/javascript",
  "globals" : { },
  "source" : "var res = false;
  var i=0;

  while (!res && i < source.effectiveRoles.length) {
    var roleId = source.effectiveRoles[i];
    if (roleId != null && roleId.indexOf("/") != -1) {
      var roleInfo = openidm.read(roleId);
      logger.warn("Role Info : {}",roleInfo);
      res = (((roleInfo.properties.name === 'Agent')
        &&(linkQualifier ==='agent'))
        || ((roleInfo.properties.name === 'Insured')
        &&(linkQualifier ==='insured')));
    }
    i++;
  }
  res"
}
```

9.2. External LDAP Configuration

Configure the LDAP server as for sample 2, (see Section 3.1.1, "LDAP Server Configuration").

The LDAP user must have write access to create users from OpenIDM on the LDAP server. When you configure the LDAP server, import the LDIF data file for this sample, *openidm/samples/multiaccountlinking/data/Example.ldif*.

9.3. Running the Multi-Account Linking Sample

This section assumes that you have installed and configured OpenDJ, as described in the previous section.

1. Prepare OpenIDM as described in Section 1.3, "Preparing the Server", then start OpenIDM with the configuration for the Multi-Account Linking sample:

```
$ cd /path/to/openidm
```

```
$ ./startup.sh -p samples/multiaccountlinking
```

2. Create the managed users.

For the purpose of this sample, create managed user entries for John Doe and Barbara Jensen. To create these users from the Admin UI, navigate to <https://localhost:8443/admin> and click Manage > User > New User.

Alternatively, use the following REST calls to create these managed user entries:


```
$ curl \
--header "Content-Type: application/json"
\
--header "X-OpenIDM-Username: openidm-admin"
\
--header "X-OpenIDM-Password: openidm-admin"
\
--request POST
\
--data '{
  "displayName": "Barbara Jensen",
  "description": "Created for OpenIDM",
  "givenName": "Barbara",
  "mail": "bjensen@example.com",
  "telephoneNumber": "1-360-229-7105",
  "sn": "Jensen",
  "userName": "bjensen",
  "accountStatus": "active"
}' \
"http://localhost:8080/openidm/managed/user?_action=create"
{
  "_id": "580f1441-ff8e-434b-9605-90e10a6fbdf6",
  "_rev": "1",
  "displayName": "Barbara Jensen",
  "description": "Created for OpenIDM",
  "givenName": "Barbara",
  "mail": "bjensen@example.com",
  "telephoneNumber": "1-360-229-7105",
  "sn": "Jensen",
  "userName": "bjensen",
  "accountStatus": "active",
  "effectiveRoles": [],
  "effectiveAssignments": []
}
```

```
$ curl \
--header "Content-Type: application/json" \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request POST \
--data '{
  "displayName": "John Doe",
  "description": "Created for OpenIDM",
  "givenName": "John",
  "mail": "jdoe@example.com",
  "telephoneNumber": "1-415-599-1100",
  "sn": "Doe",
  "userName": "jdoe",
  "accountStatus": "active"
}' \
"http://localhost:8080/openidm/managed/user?_action=create"
{
  "_id": "02632173-e413-4af1-8495-f749d5880226",
  "_rev": "1",
  "displayName": "John Doe",
  "description": "Created for OpenIDM",
  "givenName": "John",
  "mail": "jdoe@example.com",
  "telephoneNumber": "1-415-599-1100",
  "sn": "Doe",
  "userName": "jdoe",
  "accountStatus": "active",
  "effectiveRoles": [],
  "effectiveAssignments": []
}
```

In the output, you will see an `_id` associated with each user, for example, `580f1441-ff8e-434b-9605-90e10a6fbd6f6`. Make a note of these `_ids` because you will need them when you assign roles to the managed users.

3. Set up the managed roles.

This sample uses two managed roles, Agent, and Insured, to distinguish between the two user types described previously. To set up these roles in the Admin UI, navigate to <https://localhost:8443/admin> and click Manage > Role > New Role.

Alternatively, use the following REST calls to set up the two managed roles:

```
$ curl \
--header "Content-Type: application/json" \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request POST \
--data '{
  "name": "Agent",
  "description": "Role assigned to insurance agents."
}' \
"http://localhost:8080/openidm/managed/role?_action=create"
{
  "_id": "1b58ec8d-fae2-4b28-a5cf-b63567e4cf3f",
  "_rev": "1",
  "name": "Agent",
  "description": "Role assigned to insurance agents."
}
```

```
$ curl \
--header "Content-Type: application/json" \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request POST \
--data '{
  "name": "Insured",
  "description": "Role assigned to insured customers."
}' \
"http://localhost:8080/openidm/managed/role?_action=create"
{
  "_id": "617368f2-fa4e-44a2-a25a-f0a86e16ef00",
  "_rev": "1",
  "name": "Insured",
  "description": "Role assigned to insured customers."
}
```

Take note of the `_id` of each managed role. You will need these IDs to assign the roles to your managed users.

Note

Although you can use descriptive names for your role IDs by supplying them during the create operation, it is generally best practice to you use immutable server-assigned IDs.

4. Grant the managed roles to the users.

This step assumes that user `jdoh` is both an Agent and an Insured customer and that user `bjensen` is just an Insured customer of Example.com.

To grant the roles, you need the `_id` for each user and the `_id` of each role, that you obtained when you created the users and roles.

The following command grants the Agent role to jdoe:

```
$ curl \
--header "Content-type: application/json"
\
--header "X-OpenIDM-Username: openidm-admin"
\
--header "X-OpenIDM-Password: openidm-admin"
\
--request PATCH
\
--data '[
  {
    "operation" : "add",
    "field" : "/roles/-",
    "value" : {
      "_ref" : "managed/role/1b58ec8d-fae2-4b28-a5cf-b63567e4cf3f"
    }
  }
]' \
"http://localhost:8080/openidm/managed/user/02632173-e413-4af1-8495-f749d5880226"
{
  "_id": "02632173-e413-4af1-8495-f749d5880226",
  "_rev": "3",
  "displayName": "John Doe",
  "description": "Created for OpenIDM",
  "givenName": "John",
  "mail": "jdoe@example.com",
  "telephoneNumber": "1-415-599-1100",
  "sn": "Doe",
  "userName": "jdoe",
  "accountStatus": "active",
  "effectiveRoles": [
    {
      "_ref": "managed/role/1b58ec8d-fae2-4b28-a5cf-b63567e4cf3f"
    }
  ],
  "effectiveAssignments": []
}
```

The following command grants the Insured role to user bjensen:

```
$ curl \
--header "Content-type: application/json"
\
--header "X-OpenIDM-Username: openidm-admin"
\
--header "X-OpenIDM-Password: openidm-admin"
\
--request PATCH
\
--data '[
  {
    "operation" : "add",
```

```

    "field" : "/roles/-",
    "value" : {
      "_ref" : "managed/role/617368f2-fa4e-44a2-a25a-f0a86e16ef00"
    }
  }
]'\
"http://localhost:8080/openidm/managed/user/580f1441-ff8e-434b-9605-90e10a6bfd6"
{
  "_id": "580f1441-ff8e-434b-9605-90e10a6bfd6",
  "_rev": "3",
  "displayName": "Barbara Jensen",
  "description": "Created for OpenIDM",
  "givenName": "Barbara",
  "mail": "bjensen@example.com",
  "telephoneNumber": "1-360-229-7105",
  "sn": "Jensen",
  "userName": "bjensen",
  "accountStatus": "active",
  "effectiveRoles": [
    {
      "_ref": "managed/role/617368f2-fa4e-44a2-a25a-f0a86e16ef00"
    }
  ],
  "effectiveAssignments": []
}

```

The following command grants the Insured role to jdoe, because he is both an insured customer and an agent:

```

$ curl \
--header "Content-type: application/json" \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request PATCH \
--data '[
  {
    "operation" : "add",
    "field" : "/roles/-",
    "value" : {
      "_ref" : "managed/role/617368f2-fa4e-44a2-a25a-f0a86e16ef00"
    }
  }
]'\
"http://localhost:8080/openidm/managed/user/02632173-e413-4af1-8495-f749d5880226"
{
  "_id": "02632173-e413-4af1-8495-f749d5880226",
  "_rev": "4",
  "displayName": "John Doe",
  "description": "Created for OpenIDM",
  "givenName": "John",
  "mail": "jdoe@example.com",
  "telephoneNumber": "1-415-599-1100",
  "sn": "Doe",
  "userName": "jdoe",

```

```

"accountStatus": "active",
"effectiveRoles": [
  {
    "_ref": "managed/role/1b58ec8d-fae2-4b28-a5cf-b63567e4cf3f"
  },
  {
    "_ref": "managed/role/617368f2-fa4e-44a2-a25a-f0a86e16ef00"
  }
],
"effectiveAssignments": []
}

```

Note from the output that `jdoue` now has two managed roles (and thus, two values for his `effectiveRoles` property).

5. Create the managed assignments.

Assignments specify what a role actually does on a target system. A single account frequently has different functions on a system. For example, while agents might be members of the Contractor group, insured customers might be part of a Chat Users group (possibly for access to customer service). The following commands create two managed assignments that will be attached to the agent and insured roles. Note the `_id` of each assignment because you will need these when you attach the assignment to its corresponding role.

The following command creates an `ldapAgent` assignment. Users who have this assignment will have their `ldapGroups` property in OpenDJ set to `cn=Contractors,ou=Groups,dc=example,dc=com`. The assignment is associated with the `agent` link qualifier:

```

$ curl \
--header "Content-Type: application/json" \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request POST \
--data '{
  "name": "ldapAgent",
  "description": "LDAP Agent Assignment",
  "mapping": "managedUser_systemLdapAccounts",
  "attributes": [
    {
      "name": "ldapGroups",
      "value": [
        "cn=Contractors,ou=Groups,dc=example,dc=com"
      ],
      "assignmentOperation": "mergeWithTarget",
      "unassignmentOperation": "removeFromTarget"
    }
  ],
  "linkQualifiers": ["agent"]
}' \
"http://localhost:8080/openidm/managed/assignment?_action=create"
{
  "_id": "cc0dbcdc-64a4-4f5b-aade-648fc012e2b5",
  "_rev": "1",
  "name": "ldapAgent",
  "description": "LDAP Agent Assignment",
  "mapping": "managedUser_systemLdapAccounts",

```

```

"attributes": [
  {
    "name": "ldapGroups",
    "value": [
      "cn=Contractors,ou=Groups,dc=example,dc=com"
    ],
    "assignmentOperation": "mergeWithTarget",
    "unassignmentOperation": "removeFromTarget"
  }
],
"linkQualifiers": [
  "agent"
]
}

```

The following command creates an `ldapCustomer` assignment. Users who have this assignment will have their `ldapGroups` property in OpenDJ set to `cn=Chat Users,ou=Groups,dc=example,dc=com`. The assignment is associated with the `insured` link qualifier:

```

$ curl \
--header "Content-Type: application/json" \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request POST \
--data '{
  "name": "ldapCustomer",
  "description": "LDAP Customer Assignment",
  "mapping": "managedUser_systemLdapAccounts",
  "attributes": [
    {
      "name": "ldapGroups",
      "value": [
        "cn=Chat Users,ou=Groups,dc=example,dc=com"
      ],
      "assignmentOperation": "mergeWithTarget",
      "unassignmentOperation": "removeFromTarget"
    }
  ],
  "linkQualifiers": ["insured"]
}' \
"http://localhost:8080/openidm/managed/assignment?_action=create"
{
  "_id": "56b1f300-7156-4110-9b23-2052c16dd2aa",
  "_rev": "1",
  "name": "ldapCustomer",
  "description": "LDAP Customer Assignment",
  "mapping": "managedUser_systemLdapAccounts",
  "attributes": [
    {
      "name": "ldapGroups",
      "value": [
        "cn=Chat Users,ou=Groups,dc=example,dc=com"
      ],
      "assignmentOperation": "mergeWithTarget",
      "unassignmentOperation": "removeFromTarget"
    }
  ],
  "linkQualifiers": [

```

```

    "insured"
  ]
}

```

6. Add the assignments to their respective roles.

Add the `ldapCustomer` assignment to the Insured customer role:

```

$ curl \
--header "Content-type: application/json" \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request PATCH \
--data '[
  {
    "operation" : "add",
    "field" : "/assignments/-",
    "value" : {
      "_ref" : "managed/assignment/56b1f300-7156-4110-9b23-2052c16dd2aa"
    }
  }
]' \
"http://localhost:8080/openidm/managed/role/617368f2-fa4e-44a2-a25a-f0a86e16ef00"
{
  "_id": "617368f2-fa4e-44a2-a25a-f0a86e16ef00",
  "_rev": "2",
  "name": "Insured",
  "description": "Role assigned to insured customers."
}

```

7. Add the `ldapAgent` assignment to the Agent role:

```

$ curl \
--header "Content-type: application/json" \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request PATCH \
--data '[
  {
    "operation" : "add",
    "field" : "/assignments/-",
    "value" : {
      "_ref" : "managed/assignment/cc0dbcdc-64a4-4f5b-aade-648fc012e2b5"
    }
  }
]' \
"http://localhost:8080/openidm/managed/role/1b58ec8d-fae2-4b28-a5cf-b63567e4cf3f"

```


9.4. Reconciling Managed Users to the External LDAP Server

With the managed roles and assignments set up, you reconcile the managed user repository with the OpenDJ data store:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request POST \
"http://localhost:8080/openidm/recon?_action=recon&mapping=managedUser_systemLdapAccounts"
```

With the roles, assignments and link qualifiers that you set up in the previous step, this reconciliation creates three new accounts in OpenDJ:

- Two accounts under `ou=Customers,dc=example,dc=com` (one for each user who has the insured customers role), `bjensen` and `jdoe`.
- One account under `ou=Contractors,dc=example,dc=com` (for the user who has the agents role), `jdoe`.

Note that both of these users already existed in OpenDJ (from the `Example.ldif` file that you imported during the setup). If you query the list of users in OpenDJ now, you will see the multiple accounts created for `jdoe` and `bjensen` as a result of the reconciliation:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/system/ldap/account?_queryId=query-all-ids"
{
  "result": [
    {
      "_id": "uid=jdoe,ou=People,dc=example,dc=com",
      "dn": "uid=jdoe,ou=People,dc=example,dc=com"
    },
    {
      "_id": "uid=bjensen,ou=People,dc=example,dc=com",
      "dn": "uid=bjensen,ou=People,dc=example,dc=com"
    },
    {
      "_id": "uid=jdoe,ou=Contractors,dc=example,dc=com",
      "dn": "uid=jdoe,ou=Contractors,dc=example,dc=com"
    },
    {
      "_id": "uid=bjensen,ou=Customers,dc=example,dc=com",
      "dn": "uid=bjensen,ou=Customers,dc=example,dc=com"
    },
    {
      "_id": "uid=jdoe,ou=Customers,dc=example,dc=com",
      "dn": "uid=jdoe,ou=Customers,dc=example,dc=com"
    }
  ]
},
...
}
```

Chapter 10

The Trusted Servlet Filter Sample

This sample demonstrates how to use a custom servlet filter and the "Trusted Request Attribute Authentication Module" in OpenIDM. Once configured, OpenIDM can use the servlet filter to authenticate through another service.

If you want to set up authentication through OpenAM, refer to Chapter 11, *"Integrating IDM With the ForgeRock Identity Platform"*.

10.1. Before You Start

Before you start this sample, complete the following steps:

- Prepare a fresh installation of OpenIDM. (See Section 1.3, "Preparing the Server").
- Download and install the Apache Maven build tool.
- Build the custom servlet filter bundle file:

```
$ cd /path/to/openidm/samples/trustedServletfilter/filter
$ mvn clean install
```

- Copy the newly built servlet bundle file to the `openidm/bundle` directory:

```
$ cp target/sample-trusted-servletfilter-1.0.jar /path/to/openidm/bundle
```

10.2. The Sample Servlet Filter

You just built a bundle file from a Java file in the following `trustedServletfilter/filter` subdirectory: `src/main/java/org/forgerock/openidm/sample/trustedServletfilter`. The file is named `SampleTrustedServletFilter.java`.

The following line looks for the `X-Special-Trusted-User` header, to identify a specific User ID as a "trusted" user.

```
final String specialHeader = ((HttpServletRequest) servletRequest).getHeader("X-Special-Trusted-User");
```

The next line sets the special Servlet attribute `X-ForgeRock-AuthenticationId` to this trusted User ID.

```
servletRequest.setAttribute("X-ForgeRock-AuthenticationId", specialHeader);
```

The rest of the servlet filter chain continues request processing:

```
filterChain.doFilter(servletRequest, servletResponse);
```

This sample includes a `servletfilter-trust.json` file that calls the compiled OpenIDM trusted servlet `filterClass`:

```
{
  "classPathURLs" : [ ],
  "systemProperties" : { },
  "requestAttributes" : { },
  "scriptExtensions" : { },
  "initParams" : { },
  "urlPatterns" : [
    "/"*
  ],
  "filterClass" : "org.forgerock.openidm.sample.trustedervletfilter.SampleTrustedServletFilter"
}
```

10.3. Run the Sample

Start OpenIDM with the configuration for the trusted filter sample.

```
$ cd /path/to/openidm
$ ./startup.sh -p samples/trustedervletfilter
Executing ./startup.sh...
Using OPENIDM_HOME: /path/to/openidm
Using PROJECT_HOME: /path/to/openidm
Using OPENIDM_OPTS: -Xmx1024m -Xms1024m
Using LOGGING_CONFIG: -Djava.util.logging.config.file=/path/to/openidm/samples/trustedervletfilter/conf/logging.properties
Using boot properties at /path/to/openidm/samples/trustedervletfilter/conf/boot/boot.properties
-> OpenIDM ready
```

10.4. Create a Trusted User

In this section, you will create a user, and then apply the special request header `X-Special-Trusted-User` to authenticate that user.

Create user Barbara Jensen in OpenIDM, with `userName bjensen`:

```
$ curl \
--cacert self-signed.crt \
--header "X-OpenIDM-Password: openidm-admin" \
--header "X-OpenIDM-Username: openidm-admin" \
--header "Content-Type: application/json" \
--request PUT \
--data '
{
  "userName": "bjensen",
  "telephoneNumber": "6669876987",
  "givenName": "Barbara",
  "sn": "Jensen",
  "description": "Example User",
  "mail": "bjensen@example.com",
  "authzRoles" : [
    {
      "_ref" : "repo/internal/role/openidm-authorized"
    }
  ]
}' \
"https://localhost:8443/openidm/managed/user/bjensen"
```

Now you can demonstrate the servlet filter by configuring it with the noted special header. Normally, a servlet filter used for authentication does not allow a client to masquerade as any user. This sample demonstrates a basic use of a servlet filter by establishing the authentication ID.

```
$ curl \
--cacert self-signed.crt \
--header "X-Special-Trusted-User: bjensen" \
--request GET \
"https://localhost:8443/openidm/info/login?_fields=authenticationId,authorization"
```

The output should include a JSON structure with the user's authentication and authorization details. In this case, user **bjensen** is authenticated with the "openidm-authorized" role.

```
{
  "_id" : "",
  "authenticationId" : "bjensen",
  "authorization" : {
    "id" : "bjensen",
    "component" : "managed/user",
    "roles" : [ "openidm-authorized" ]
  }
}
```

10.5. Customizing the Sample for an External System

To customize this sample for an external authentication/authorization system, you need a servlet filter which authenticates against that external system. You may use a third-party supplied filter, or develop your own filter, using the one in this sample as a model.

The filter you use should have at least the following capabilities:

- Perform REST calls to another system.
- Search through databases.
- Inspect headers related to authentication and authorization requests.

This servlet filter must set the username of the authenticated user in a special request attribute. You need to configure that same attribute name in the `TRUSTED_ATTRIBUTE` authentication module, specifically the value of `authenticationIdAttribute`.

It is helpful if you have a filter that returns an object with the `userRoles` property. If your filter does not support queries using the following parameter:

```
queryOnResource + "/" + authenticationId
```

You will need to provide a security context augmentation script that populates the following authorization properties in the "security" object:

- `security.authorization.component`
- `security.authorization.roles`

The value for the `security.authorization.component` is automatically set to the value specified in any existing `queryOnResource` property.

Chapter 11

Integrating IDM With the ForgeRock Identity Platform

This sample demonstrates the integration of three products within the ForgeRock Identity Platform. The sample shows ForgeRock Access Management, (AM), ForgeRock Directory Services (DS), and ForgeRock Identity Management (IDM) working together to manage identities and authentication.

In this sample, you'll set up an *OpenID Connect Authorization Code Flow*, where OpenIDM acts as the Relying Party, and OpenAM takes the role of the OpenID Provider.

After you've demonstrated integration between the noted three ForgeRock products, you'll see how you can integrate customer identity and access management, in Section 11.10, "Integrating Social ID Logins".

11.1. Preparing Your Systems

In this section, you'll prepare dedicated systems for the installation of OpenIDM, OpenAM, and OpenDJ.

Note

This sample assumes that you will install ForgeRock Identity Management 5, ForgeRock Access Management 5, and ForgeRock Directory Services 5.

OpenAM requires the use of a Fully-Qualified Domain Name (FQDN). For consistency, this sample includes FQDNs for OpenIDM, OpenAM and, OpenDJ on either an appropriate DNS server or the `hosts` file for each system.

This sample assumes that you have assigned the following FQDNs to the OpenAM, OpenDJ, and OpenIDM systems, respectively:

- `openam.example.com`
- `opendj.example.com`
- `openidm.example.com`

Note

The `example.com` domain is used in this sample, for consistency with the `Example.ldif` file that you'll find in the `openidm/samples/fullStack/data` directory. If you want to use a different domain with this sample, change the data in the `Example.ldif` file accordingly.

Refer to the following sections of each product document to prepare your systems.

- For OpenIDM, see Section 1.1, "Before You Install" in the *Installation Guide*.
- For OpenAM, see *Preparing for Installation* in the *Access Management Installation Guide*.
- For OpenDJ, see *Chapter 1, Before You Install* in the *Directory Services Installation Guide*.

11.2. Preparing an Instance of OpenDJ

Configure the OpenDJ server as described in Section 3.1.1, "LDAP Server Configuration".

You need to configure the OpenDJ server to support write access. This allows you to create users from OpenIDM or OpenAM on the same LDAP server. When you configure the LDAP server, import the LDIF file associated with the Full Stack sample: (`openidm/samples/fullStack/data/Example.ldif`).

When you configure OpenAM, use the following information to configure OpenDJ as an external data store:

- Access URL and port for the LDAP server; for this sample, we use `opendj.example.com:1389`.
- LDAP Bind DN, normally `cn=Directory Manager`.
- LDAP Bind Password, which should match the password configured the LDAP server.

11.3. Starting IDM

Prepare OpenIDM, as described in Section 1.3, "Preparing the Server", then start OpenIDM with the configuration for the full stack sample.

```
$ cd /path/to/openidm
$ ./startup.sh -p samples/fullStack
```

11.4. Installing OpenAM for Integration

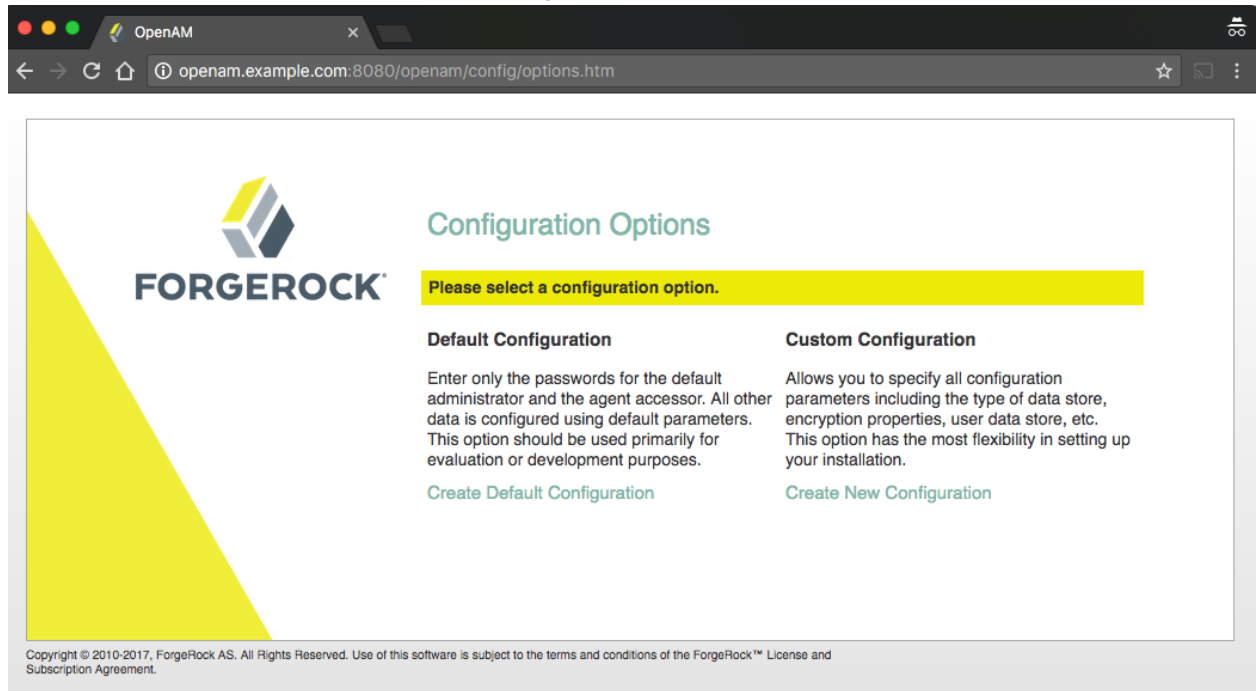
This section is based on the following chapter in the Access Management Installation Guide, *To Custom Configure an Instance*.

In that chapter, you'll deploy the OpenAM `.war` file to the appropriate web application container directory. OpenAM supports multiple web application containers. For one example, see *Deploying in the Access Management Installation Guide*.

Note

This sample uses the default unencrypted web container port (8080) for convenience. In production, you should set up OpenAM over a secure port, as described in the following chapter of the Access Management Installation Guide: *Securing Installations*.

If you followed the procedure associated with the OpenAM `.war` file, and configured the FQDNs described in Section 11.1, "Preparing Your Systems", point your browser to <http://openam.example.com:8080/openam>. You should see the following screen:



Procedure 11.1. Installing OpenAM for This Sample

To install OpenAM, read through the following procedure: *To Custom Configure an Instance in the Access Management Installation Guide*.

Follow the first steps of the noted procedure.

When you install OpenAM for this sample, include the settings described in the following steps:

1. When you see the **Server Settings** step, verify that the server settings are valid for your configuration. Make sure these settings correspond to the FQDN for your instance of OpenAM. For this sample, you'd specify:

- **Server URL:** `http://openam.example.com:8080`
- **Cookie Domain:** `openam.example.com`

2. In the Configuration Data Store Settings screen, accept all but one of the defaults to allow OpenAM to store configuration data in an embedded directory.

This sample assumes that this is your first and only instance of OpenAM, which is the default. Include a Root Suffix that matches the root domain for your instance of OpenAM, in this case, `dc=example,dc=com`.

OpenAM Configurator

Custom Configuration Option

1. General
2. Server Settings
→ **Configuration Store**
4. User Store
5. Site Configuration
6. Agent Information
7. Summary

Step 3: Configuration Data Store Settings

If no other OpenAM instance already exists in the environment, then choose First Instance. If one or more OpenAM instances already exist in the environment, choose Add to Existing Deployment.

First Instance Add to Existing Deployment?

* Indicates required field

Configuration Store Details

Configuration Data Store OpenAM OpenDJ

* SSL/TLS Enabled

* Host Name

* Port

* Admin Port

* JMX Port

* Encryption Key

* Root Suffix OK

Previous **Next** **Cancel**

Note

If your configuration of OpenDJ includes a different Root Suffix, modify your OpenAM installation settings accordingly.

- In the User Data Store Settings screen, configure where OpenAM looks for user identities. In this case, select **Other User Data Store** and choose **OpenDJ** as the User Data Store type. You can then point this installation to the OpenDJ server configured at `opendj.example.com`, on port 1389, with a root suffix of `dc=example,dc=com`.

OpenAM Configurator

Custom Configuration Option

1. General
2. Server Settings
3. Configuration Store
→ 4. **User Store**
5. Site Configuration
6. Agent Information
7. Summary

Step 4: User Data Store Settings

You can use the data store that comes with the OpenAM configuration data store, or you can use a different user data store. A good practice for setting up production environments is to use an external user data store, one that is different than the OpenAM user data store. Please note that Policy Service and LDAP Authentication Module shall be configured to use the Directory Administrator DN and Password provided here.

OpenAM User Data Store
 Other User Data Store

* Indicates required field

User Store Details

* User Data Store Type OpenDJ Oracle Directory Server Enterprise Edition
 AD with Domain Name Active Directory with Host and Port
 IBM Tivoli Directory Server Active Directory Application Mode

* SSL/TLS Enabled

* Directory Name

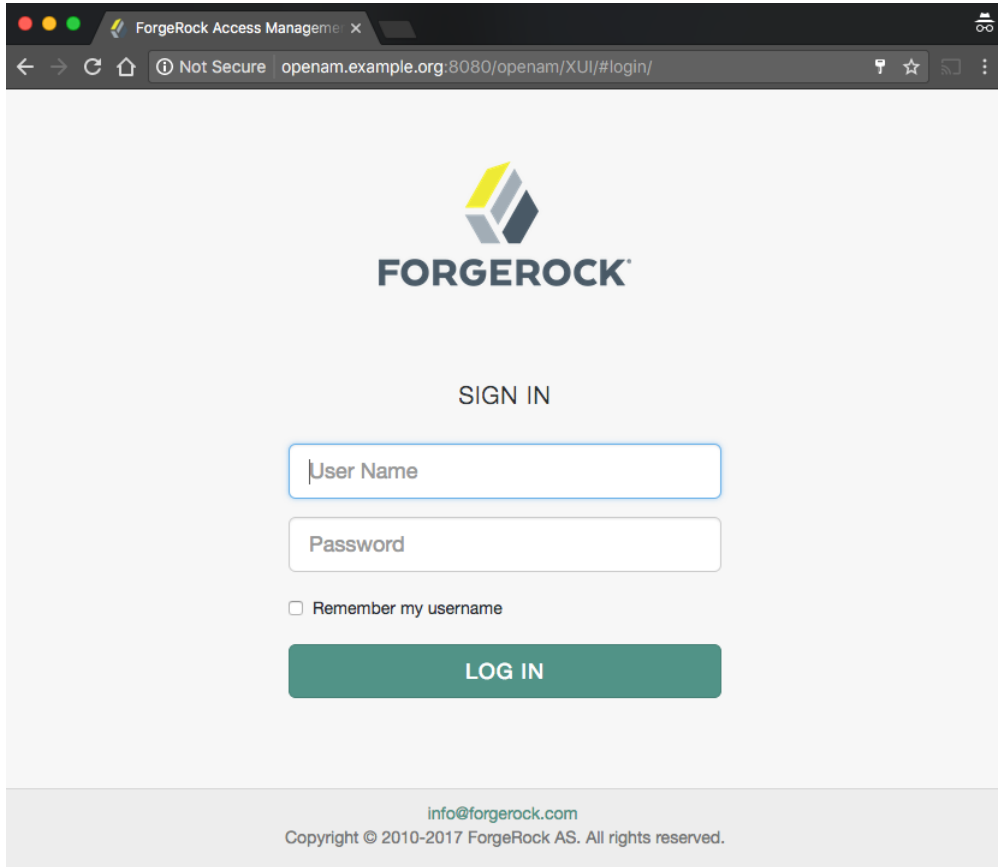
* Port OK

* Root Suffix OK

* Login ID OK

* Password OK

- For this sample, you're installing one instance of OpenAM. So when you reach the Site Configuration screen, accept the default **No**, and select **Next** to continue.
- Complete the AM installation process as directed, until you're logged into AM. For details, see the previously cited Access Management Installation Guide procedure: *To Custom Configure an Instance*



Now proceed with the configuration of OpenAM as an OpenID Connect provider for OpenIDM, as described in the following section:

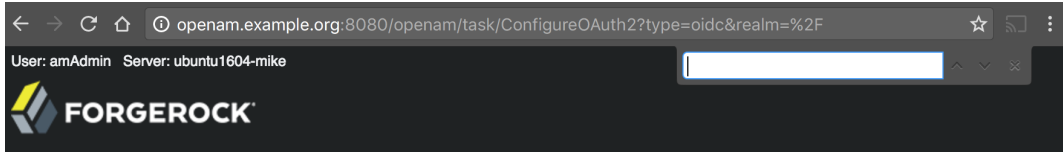
11.5. Configuring OpenAM for Integration with IDM

To integrate OpenAM and OpenIDM, you need to configure OpenAM as an OpenID Connect provider. OpenID Connect 1.0 is an authentication service built on OAuth 2.0. In the following sections, you'll configure OpenID Connect 1.0 and OAuth 2.0 settings, using information from OpenIDM.

11.5.1. Configuring OpenAM as an OAuth 2.0 Provider for IDM

In this section, you'll set up OpenAM as an OAuth 2.0 provider. To do so, take the following steps:

1. Navigate to <http://openam.example.com:8080/openam>. Log into OpenAM as an administrator, with username `amadmin`, with the password you created during Procedure 11.1, "Installing OpenAM for This Sample".
2. For the purpose of this sample, choose the default Top Level Realm.
3. Select Configure OAuth Provider.
4. Select Configure OpenID Connect. You should see the following window:



Configure OpenID Connect

[Create](#) [Cancel](#)

Configure OpenAM as an OpenID Connect authorization server. The provider service will be configured with settings that conform to the [OpenID Connect specification](#), which you can modify if required.

* Indicates required field

* Realm:

Configure OAuth2/OpenID Connect Service

* Refresh Token Lifetime (seconds):
The time in seconds a refresh token is valid for

* Authorization Code Lifetime (seconds):
The time in seconds an authorization code is valid for

* Access Token Lifetime (seconds):
The time in seconds an access token is valid for

Issue Refresh Tokens:
Check to enable generation of refresh tokens

Issue Refresh Tokens on Refreshing Access Tokens:
Check to enable generation of refresh tokens when refreshing access tokens

* Scope Implementation Class:
The class that contains the required scope implementation, must implement the `org.forgerock.oauth2.core.ScopeValidator` interface

Configure OAuth2 Authorization End Point Protection Policy

A policy to protect the OAuth2 authorization end point will be created. This policy will be named OAuth2ProviderPolicy. The policy will protect the endpoint <http://openam.server.name.com/openam/oauth2/authorize>. The purpose of this policy is to redirect clients to the OpenAM login page to authenticate a resource owner each time they go to the authorize end point. To do advanced policy management can be done using the policies tab for each realm.

Register OAuth2/OpenID Connect Client(s)

The last step is to register client(s) for the OAuth2/OpenID Connect Provider to issue tokens to. Clients can be registered by navigating to the OpenAM agents tab and selecting OAuth 2.0/OpenID Connect Client. A Client can also be registered by visiting the jsp registration page at <http://openam.server.name.com/openam/oauth2/registerClient.jsp>

5. Accept the defaults. Choose the Create button in the Configure OpenID Connect window.
6. You'll see a window with the following message:

```
OAuth2 Configured for the given realm
```

```
Successfully configured OAuth2 for realm /.
```

```
A policy was created on the realm for the authorization end point. The policy name  
is: OAuth2ProviderPolicy
```

Select OK.

For detailed information on OpenAM's OAuth 2.0 Provider configuration, see the *Access Management OAuth 2.0 Guide*.

11.5.2. Getting Information From IDM

OpenIDM includes helpful information for the next step in the OpenAM configuration process, Section 11.5.3, "Configuring OpenAM's OpenID Connect Agent". To see this information from the Admin UI, navigate to **Configure > Authentication**. In the window that appears, select **ForgeRock Identity Provider**. Note the values for the **End User Redirect URI** and the **Admin Redirect URI**.

Note

Make a copy of the information in the following figure. Do *not* press Submit until you've gone through Section 11.6, "Plugging IDM Into OpenAM".

Figure 11.1. Use The Information In This Figure

Configure ForgeRock Identity Provider ×

1. Configure Access Manager

Configure the OpenAM OpenID Connect Client Agent with these redirection URIs. Visit the **OpenID Connect 1.0 Guide** for more information.

End User Redirect URI

Admin Redirect URI

2. Enter configuration details

Well-Known Endpoint
For more information visit the [OpenID Connect 1.0 Guide](#)

Client ID

Client Secret

Route to OpenAM User Datastore

For an explanation of these values, see Table 11.1, "Configure the OpenAM OAuth 2.0/OpenID Connect Client With the Following Options".

11.5.3. Configuring OpenAM's OpenID Connect Agent

In this section, you'll configure an OpenAM Client Agent that allows you to register your instance of OpenIDM as an OpenID Connect 1.0 Client.

To proceed, navigate to the OpenAM UI at <http://openam.example.com:8080/openam>. Log into the UI as the OpenAM administrative user, `amadmin`, and the password you created for that user earlier in this chapter.

Select Top Level Realm > Applications > Agents. In the window that appears, select Agents > OAuth 2.0/OpenID Connect Client.

In the Agent sub-window, select New. In the New OAuth 2.0/OpenID Connect Client window that appears, you'll assign a name and password that corresponds to the `client_id` and `client_secret` properties associated with the *OpenID Connect Core 1.0* specification.

New OAuth 2.0/OpenID Connect Client

* **Name:**

* **Password:**

* **Re-Enter Password:**

For this sample, set a name (`client_id`) of `openidm` and a password (`client_secret`) of `openidm`.

Choose Create to create a new agent. This returns you to the OAuth 2.0/OpenID Connect Client tab, where you can now select your newly created `openidm` agent. Do so.

General	Authentication	Services	Data Stores	Privileges	Policies	Subjects	Agents
Web	J2EE	2.2 Agents	OAuth 2.0/OpenID Connect Client	Agent Authenticator	SOAP STS Agent		

/ (Top Level Realm)

OAuth 2.0/OpenID Connect Client
OAuth 2.0 clients use access tokens issued by OpenAM to access protected resources.

Agent (1 Agent(s))	
<input type="button" value="New..."/>	<input type="button" value="Delete"/> <input type="button" value="Refresh"/>
<input checked="" type="checkbox"/> <input type="button" value="Edit"/>	Name
<input type="checkbox"/>	openidm
	<input type="button" value="Refresh"/>

Configure the `openidm` client agent. Make changes based on the properties shown in the following table:

Table 11.1. Configure the OpenAM OAuth 2.0/ OpenID Connect Client With the Following Options

Property Name	Description	Entry
Redirection URIs (IDM End User Redirect URI)	Supports URIs with hash fragments; add to the Redirection URIs text box	http://openidm.example.com:8080/oauthReturn.html
Redirection URIs (IDM Admin Redirect URI)	Supports URIs with hash fragments; add to the Redirection URIs text box	http://openidm.example.com:8080/admin/oauthReturn.html
Scope	Must include <code>openid</code> as a supported scope to use the <code>/oauth2/.well-known/openid-configuration</code> endpoint. For more information, see <i>Configuring for OpenID Connect Discovery from the Access Management OpenID Connect 1.0 Guide</i>	<code>openid</code>
Token Endpoint Authentication Method	Supports secure authentication through OpenAM	<code>client_secret_post</code>
Post Logout Redirect URI	Specify the URI to which to redirect the user-agent after the client logout process	See Limits on the Post Logout Redirect URI
Implied Consent	Bypasses any consent requirements of the resource owner	Enabled

For a description of each of the property values shown on the `openidm` client agent page, see *OAuth 2.0 and OpenID Connect 1.0 Client Settings*, from the *Access Management OpenID Connect 1.0 Guide*.

Limits on the Post Logout Redirect URI

The Post Logout Redirect URIs for this sample must match the URIs expected for OpenIDM. In this case, for the Admin and Self-Service UIs, you'd include the following values:

```
http://openidm.example.com:8080/admin/
http://openidm.example.com:8080/
```

If you've configured connections over a secure port, modify the Post Logout Redirect URIs accordingly.

Save your changes. Scroll to the top of the screen, and select Save.

11.5.4. Integrating OpenAM's OAuth2 Provider Service

Now that you've configured OpenAM's OpenID Connect Agent, you can configure OpenAM's OAuth2 Provider Service. To do so in OpenAM, take the following steps:

1. Navigate to <http://openam.example.com:8080/openam/>. Log into OpenAM as an administrator, with username `amadmin`, and the password you created during Procedure 11.1, "Installing OpenAM for This Sample".
2. Choose the default Top Level Realm.
3. Choose Services > OAuth2 Provider.
4. Navigate to the Advanced tab, and enable *Allow Clients to Skip Consent*; for more information, see the Access Management OAuth 2.0 Guide on *Allowing Clients to Skip Consent*.

Select *Save Changes*.

11.6. Plugging IDM Into OpenAM

Now that you've configured an OpenID Connect Client on OpenAM, you can configure OpenIDM as an OpenID Connect *Relying Party* to authenticate through OpenAM. Return to the OpenIDM Admin UI. Navigate to Configure > Authentication. In the window that appears, select ForgeRock Identity Provider.

Configure ForgeRock Identity Provider ✕

1. Configure Access Manager

Configure the OpenAM OpenID Connect Client Agent with these redirection URIs. Visit the [OpenID Connect 1.0 Guide](#) for more information.

End User Redirect URI

Admin Redirect URI

2. Enter configuration details

Well-Known Endpoint

For more information visit the [OpenID Connect 1.0 Guide](#)

Client ID

Client Secret

Route to OpenAM User Datastore

The ForgeRock Identity Provider is also known as OpenAM.

You'll need to modify three things before the UI allows you to Submit your changes:

- **Well-Known Endpoint**

An OpenAM endpoint with related configuration information.

For more information on the OpenAM Well-Known Endpoint, see the Access Management OpenID Connect 1.0 Guide on *Configuring for OpenID Connect Discovery*.

- **Client ID**

Use the name for the OAuth 2.0/OpenID Connect Client agent that you configured in Section 11.5.3, "Configuring OpenAM's OpenID Connect Agent".

- **Client Secret**

Use the password for the OAuth 2.0/OpenID Connect Client agent that you configured in Section 11.5.3, "Configuring OpenAM's OpenID Connect Agent".

Note

You'll see one additional option: *Route to AM User Datastore*. For this sample, the default should match the repository for the external data store, in this case, `system/ldap/account`.

Choose Submit to save your changes. If successful, you'll see the following message:

Your current session may be invalid. Click here to logout and re-authenticate.

When you select the link IDM takes you to the AM Login Screen. You should now be able to connect as the AM Administrative user, `amadmin`, with the password you set when installing AM.

Note

If you're more familiar with AM, be careful. Use `amadmin` to log into the full stack. (While you can normally use `amAdmin` to log into IDM, that account name will *not* work here.)

In this configuration, when you log into AM, you're taken to the OpenIDM Admin UI.

11.6.1. Changes to Session and Authentication Modules

When you activate the ForgeRock Identity Provider, OpenIDM implements the following changes to the session module, as shown in the `authentication.json` file:

- The maximum token lifetime of the `JWT_SESSION` is shortened from the OpenIDM default of 120 minutes to 5 seconds.
- The token idle lifetime is reduced from the OpenIDM default of 30 minutes to 5 seconds.

Reducing the OpenIDM-specific session token times in this way supports a rapid reaction when the OpenAM session expires. In essence, this change means that sessions are governed by OpenAM. For more information, see the *About Sessions* section of the *Access Management Authentication and Single Sign-On Guide*.

Table 11.2. Authentication Modules When Enabling the ForgeRock Identity Provider

Authentication Module	Status
STATIC_USER	No Change
MANAGED_USER	Disabled
INTERNAL_USER	No Change
CLIENT_CERT	Disabled
PASSTHROUGH	Disabled
OPENID_CONNECT	Enabled

For more information on each module, see Section 18.1.2, "Supported Authentication and Session Modules" in the *Integrator's Guide*.

11.7. Demonstrating Integration

Once this process is complete, navigate to the OpenIDM Admin UI. For this sample, navigate to <http://openidm.example.com:8080/admin>. The settings you changed in Section 11.6, "Plugging IDM Into OpenAM", along with what you did to configure OpenAM in this sample, will redirect you to <http://openam.example.com:8080/openamLongRedirectURI>.

When you log into OpenAM, with this *longRedirectURI*, as the OpenAM administrative user (*amadmin*), it connects you to the OpenIDM Admin UI, as the OpenIDM administrative user. To understand why, read:

- Section 11.7.1, "Analyzing the OpenAM Redirect URI"
- Section 11.7.2, "Understanding the Integrated OpenAM Administrative User"

Note

If you're more familiar with AM, be careful. Use *amadmin* to log into the full stack. (While you can normally use *amAdmin* to log into IDM, that account name will *not* work here.)

11.7.1. Analyzing the OpenAM Redirect URI

To understand what happened when you configured OpenAM and OpenIDM, it may help to analyze the Redirect URI. When you do, you'll also discover other OpenID Connect-compliant information.

Based on the configuration described in this sample, the full redirect URI is as follows:

```
http://openam.example.com:8080/openam/XUI/?realm=%2F&goto=
http%3A%2F%2Fopenam.example.com%3A8080%2Fopenam%2Foauth2%2Fauthorize
%3Fresponse_type%3Dcode%26scope%3Dopenid
%26redirect_uri%3Dhttp%253A%252F%252Fopenidm.example.com%253A8080%252Fadmin%252FoauthReturn.html
%26state%3Dlogin%2526provider%253DOPENAM
%2526redirect_uri%253Dhttp%253A%252F%252Fopenidm.example.com%253A8080%252Fadmin%252FoauthReturn.html
%2526gotoURL%253D%2523%26nonce%3D1cp3yw64ag0v%26client_id%3Dopenidm#login/
```

If you analyze the full redirect URI, you should see several settings that you configured in OpenAM in this sample:

- **oauth2**: See Section 11.5.1, "Configuring OpenAM as an OAuth 2.0 Provider for IDM"
- **scope**: See Table 11.1, "Configure the OpenAM OAuth 2.0/OpenID Connect Client With the Following Options"
- **redirect_uri**: See Table 11.1, "Configure the OpenAM OAuth 2.0/OpenID Connect Client With the Following Options"

- **nonce**: See *OpenID Connect Core 1.0 Specification, Nonce Implementation Notes*
- **client_id**: See Section 11.5.3, "Configuring OpenAM's OpenID Connect Agent"

11.7.2. Understanding the Integrated OpenAM Administrative User

If you've tried the integrated OpenAM/OpenIDM system, as described in Section 11.7, "Demonstrating Integration", you'll realize that you used the OpenAM Administrative account `amadmin` to log into OpenIDM. After logging in, you can confirm that the account in use to administer OpenIDM is the default, `openidm-admin`.

To understand how that works, examine the `amSessionCheck.js` file in the following directory: `/path/to/openidm/bin/defaults/script/auth`. See the following code block, and how it gives user `amadmin` the OpenIDM administrative user `id`, along with standard administrative access.

```
if (security.authenticationId === "amadmin") {
  security.authorization = {
    "id" : "openidm-admin",
    "component" : "repo/internal/user",
    "roles" : ["openidm-admin", "openidm-authorized"],
    "moduleId" : security.authorization.moduleId
  };
}
```

You can modify this script as desired. For example, you could limit the OpenIDM roles given to the OpenAM administrative account. Alternatively, you could set up such authorization mapping to a different OpenAM account.

11.8. The Integrated `authentication.json` File

When you configure this sample as described in this chapter, OpenIDM adds the following lines to the `authentication.json` file in your project's `conf/` subdirectory.

The following excerpt includes information that appears to go beyond what you configured. OpenIDM reads the well-known endpoint that you included in Section 11.6, "Plugging IDM Into OpenAM", and includes the following information in the `authentication.json` file:

- **authorization_endpoint** as described in the *OpenID Connect Discovery 1.0* specification.
- **token_endpoint** as described in the *OpenID Connect Discovery 1.0* specification.
- **end_session_endpoint** as described in the *OpenID Connect Session Management 1.0* specification.

```
{
  "enabled" : true,
  "properties" : {
    "resolvers" : [
      {
```

```

        "name" : "OPENAM",
        "icon" : "<button class=\|\"btn btn-lg btn-default btn-block btn-social-provider
        \><img src=\|\"images/forgerock_logo.png\|\">Sign In</button>\",
        "scope" : [
            "openid"
        ],
        "well-known" : "http://openam.example.com:8080/openam/oauth2/.well-known/
        openid-configuration",
        "client_id" : "openid",
        "client_secret" : {
            "$crypto" : {
                "type" : "x-simple-encryption",
                "value" : {
                    "cipher" : "AES/CBC/PKCS5Padding",
                    "salt" : "y6e/pXQ2FHZmca10GbIW4Q==",
                    "data" : "Tp5UQwIsnP40USRepFVUkw==",
                    "iv" : "pbmXr+qyJ6/Ef0MttviSrw==",
                    "key" : "openidm-sym-default",
                    "mac" : "FDEPyfDjyGi6z0r3hKGF4g=="
                }
            }
        },
        "authorization_endpoint" : "http://openam.example.com:8080/openam/oauth2/authorize",
        "token_endpoint" : "http://openam.example.com:8080/openam/oauth2/access_token",
        "end_session_endpoint" : "http://openam.example.com:8080/openam/oauth2/connect/endSession"
    }
},
"queryOnResource" : "managed/user",
"defaultUserRoles" : [
    "openidm-authorized"
],
"openIdConnectHeader" : "authToken",
"propertyMapping" : {
    "authenticationId" : "userName",
    "userRoles" : "authzRoles"
},
"augmentSecurityContext" : {
    "type" : "text/javascript",
    "globals" : {
        "sessionValidationBaseEndpoint" : "http://openam.example.com:8080/openam/json/sessions/"
    },
    "file" : "auth/amSessionCheck.js"
}
},
"name" : "OPENID_CONNECT"
}

```

You can also configure an `OPENID_CONNECT` module without using the ForgeRock Identity Provider. To do so from the Admin UI, select `Configure > Authentication`, select the `Modules` tab, and select the `OPENID_CONNECT` module. For more information on what that adds to the `authentication.json` file, see Section 18.1.4, "Configuring Authentication With OpenID Connect" in the *Integrator's Guide*.

11.9. Reconciliation and OpenAM

In this section, you'll see how reconciliation in OpenIDM affects users seen in OpenAM.

When you first integrated the three products (OpenIDM, OpenAM, and OpenDJ), you included information from an `Example.ldif` file from the `openidm/samples/fullStack/data/` subdirectory.

When you set up OpenDJ, you populated the directory with user data from from the `Example.ldif` file in the `openidm/samples/fullStack/data/` directory. To get that user data from OpenDJ into the OpenIDM managed user datastore, run reconciliation. To do so, select Configure > Mappings. You'll see two mappings:

- `systemLdapAccounts_managedUser`, which maps information from the LDAP datastore (OpenDJ - `systemLdapAccounts`) to the OpenIDM managed user database.
- `managedUser_systemLdapAccounts`, which maps information in the reverse direction.

Before running reconciliation, open the OpenAM administrative interface. Navigate to `http://openam.example.com:8080/openam`. Log in as the OpenAM administrative user, `amadmin`, with the password you used when installing OpenAM earlier in this sample, in Section 11.4, "Installing OpenAM for Integration".

Select Top Level Realm > Subjects. You should see four accounts:

- Two default OpenAM accounts: `amadmin` and `anonymous`.
- Two accounts from the aforementioned `Example.ldif` file, which have been loaded into the common user datastore, OpenDJ.

Now return to the OpenIDM Admin UI. Select Configure > Mappings. In the `systemLdapAccounts_managedUser` mapping, choose Reconcile.

Once reconciliation is complete, you'll see users from `Example.ldif` when you select Manage > User.

Create a new OpenIDM user as described in Procedure 4.3, "To Add a User Account" in the *Integrator's Guide*. Add a password to that account.

Return to the OpenAM administrative interface at `http://openam.example.com:8080/openam`. Select Top-Level Realm > Subjects. You should see a list of users that include the user that you created in OpenIDM.

11.10. Integrating Social ID Logins

To take advantage of the authorization features provided by OpenAM, you can extend the features shown in this sample. This section demonstrates how you can integrate social identities, specifically Google and Facebook. It also assumes that you've configured the full stack sample, as described so far in this chapter.

Before you start this process, you may want to go through the social authentication process for both OpenIDM and OpenAM, as described in the following chapters:

- Section 10.3, "Setting Up Google as a Social Identity Provider" in the *Integrator's Guide*.
- Section 10.5, "Setting Up Facebook as a Social Identity Provider" in the *Integrator's Guide*.

- *Implementing Social Authentication in the Access Management Authentication and Single Sign-On Guide.*

As this setup uses OpenAM to manage authentication for social identities, you can ignore the associated OpenIDM module, described in the following section: Section 10.7, "Configuring the Social Providers Authentication Module" in the *Integrator's Guide*.

To integrate OpenAM's social provider features, you'll need to change defaults for Google and Facebook as described in the following section of the Access Management Authentication and Single Sign-On Guide: *Configuring Additional Settings for a ForgeRock Identity Platform Deployment*.

For more information on these properties, see the following section of the Access Management Authentication and Single Sign-On Guide: *OAuth 2.0/OpenID Connect Authentication Module Properties*.

11.11. Registering Users via OpenIDM

You can now register users by their social IDs through OpenIDM. To do so, take the following steps:

1. Navigate to <http://openidm.example.com:8080/#register/>. Be sure to enter the full URL, including the final forward slash `/`.
2. In the *Register Your Account* screen that appears, you should see the following options:
 - *Register with Google*
 - *Register with Facebook*
3. Select one of these options, and proceed as prompted with either a Google or Facebook account.

Review the list of users in both OpenIDM and OpenAM. You should see the same social ID accounts in both systems. Since OpenAM and OpenIDM both use the same instance of OpenDJ, the lists of users (beyond the defaults for each system) should be identical.

- In the OpenIDM Admin UI, select Manage > User.
- In the OpenAM Administrative Interface, select Realm > Subjects.

Note

If you see the following error:

```
User requires profile to login
```

Follow the procedure at the beginning of this section. You need to register that user first.

Once you've registered users, you can navigate to <http://openidm.example.com:8080/>. The configuration of this sample means that your browser gets redirected to <http://openam.example.com:8080/openam/XUI/?realm=<longURI>>.

You'll see icons for the social ID providers that you configured, such as Google and Facebook. When you select one of these icons, what happens next depends:

- If you have cleared the cache on the browser, you'll be prompted to log into your selected social ID provider.
- If you have not cleared the cache, the full stack sample assumes that you want to log in with the social ID provider account that you just configured.

The result is the same; you're taken to the OpenIDM end-user dashboard.

Chapter 12

Workflow Samples

This chapter walks you through the workflow provisioning sample (in the `openidm/samples/workflow` directory) and the workflow use cases (in the `openidm/samples/usecase` directories). For a complete list of the samples provided with OpenIDM, and an overview of each sample, see Chapter 1, "Overview of the Samples".

12.1. Sample Workflow - Provisioning User Accounts

This sample, provided in `openidm/samples/workflow`, demonstrates a typical use case of a workflow — provisioning new users.

The sample demonstrates the use of the Admin UI, to configure user self-service and the Self-Service UI that enables users to complete their registration process.

This sample simulates the following scenario:

- An existing employee requests that an outside contractor be granted access to an organization's system.
- The *system* in this case, is OpenIDM's managed user repository and a remote datasource, represented by an XML file.
- User roles are stored separately, in a CSV file.

The sample has three mappings — two for the bidirectional synchronization of the managed user repository and the XML data store, and one for the synchronization of the roles data (stored in the CSV file) to the managed repository.

12.1.1. Prepare OpenIDM For the Provisioning Sample

In this section, you start OpenIDM, configure the outbound email service, and reconcile user and role data. The reconciliation operations create two managed users, `user1` and `manager1`, and two managed roles, `employee` (assigned to `user1`) and `manager` (assigned to `manager1`).

1. Start OpenIDM with the configuration for the provisioning sample.

```
$ cd /path/to/openidm
$ ./startup.sh -p samples/workflow
```

2. Log in to the Admin UI (<https://localhost:8443/admin>) with the default username (`openidm-admin`) and password (`openidm-admin`).

3. Configure the outbound email service.

An email configuration is required to send a password-reset email to the user, at the end of the sample.

Select Configure > System Preferences > Email and provide connection information for your preferred email host.

Gmail is configured by default but you must still supply credentials to use Gmail.

4. Reconcile the role data, and the user data.

a. Select Configure > Mappings.

b. Click the first mapping (`systemRolesFileRole_managedRole`) and click Reconcile Now.

This reconciliation operation creates two roles in the managed repository (`employee` and `manager`). You can check the result of the reconciliation by selecting Manage > Role.

c. Go back to the Mappings page (Configure > Mappings), select the second mapping (`systemXmlfileAccounts_managedUser`) and click Reconcile Now.

This reconciliation operation creates the top-level managers (users who do not have their own `manager` property) in the managed user repository. In this sample, there is only one top-level manager (`manager1`).

d. Click Reconcile Now a second time.

This reconciliation operation creates the employees of the managers that were created by the previous reconciliation. In this sample, there is only one employee (`employee1`).

Check that the manager and employee entries were created correctly by selecting Manage > User. You should have two users — `manager1` and `user1`.

5. Verify the relationships between your new user and role objects:

a. Click on `user1`. Note that the `Manager` field shows `manager1` for this user.

b. On the Authorization Roles tab, note that `user1` has two roles — `employee` and `openid-authorized`.

c. Click Back to User List then click on `manager1`. Note that the `Manager` field is empty for this user.

d. On the Authorization Roles tab, note that `manager1` has two roles — `manager` and `openid-authorized`.

6. Verify the available workflows:

a. Select Manage > Processes > Definitions.

b. Click the Contractor onboarding process and look at the sequence diagram for this workflow.

7. Log out of the Admin UI by selecting Log Out from the top right dropdown menu.

OpenIDM is now prepared for the Provisioning sample. In the next section you will walk through the workflow whose sequence diagram you saw in the previous step.

12.1.2. Running the Provisioning Sample

As part of provisioning, employees are required to initiate a *Contractor Onboarding* process. This process is a request to add a contractor to the managed user repository, with an option to include the contractor in the original data source (the XML file).

When the employee has completed the required form, the request is sent to the manager for approval. Any user with the role `manager` can claim the approval task. If the request is approved, an email is sent to the address provided in the initial form, with a request for the contractor to reset their password. When the password reset has been completed, the contractor is created in the managed user repository. If a request was made to add the contractor to the original data source (the XML file) this is done when the manager approves the request.

1. Log in to the Self-Service UI (<https://localhost:8443/>) as the user you created in the previous section (`user1`), with password `Welcome1`.

The user Dashboard is displayed:

The screenshot shows a user dashboard with the following sections:

- MY TASKS**: No tasks assigned.
- MY GROUP'S TASKS**: Your group has no tasks at this time.
- PROCESSES**: A single process is listed: "Contractor onboarding process" with a green checkmark icon and a "Details" link.
- NOTIFICATIONS**: No notifications.

2. Initiate the provisioning workflow as user1:
 - a. Under Processes, click the Details link next to the Contractor onboarding process and complete the form for the sample user you will be creating.

Use an email address that you have access to because you will need the email that is sent to this address to complete the workflow.

In the Provision to XML field, select Yes.

This selection enables implicit synchronization from the managed user repository to the XML file.

✓ Contractor onboarding process ▼ Details

Contractor Details

Username	<input type="text" value="johnb"/>
Email address	<input type="text" value="johnb@example.com"/>
First Name	<input type="text" value="John"/>
Last Name	<input type="text" value="Brand"/>
Mobile Phone	<input type="text" value="14155991100"/>
Provision to XML	<input type="text" value="Yes"/>
Department	<input type="text" value="Payroll"/>
Job Title	<input type="text" value="Payroll Clerk"/>
Description	<input type="text" value="Temp payroll clerk"/>
Start Date	<input type="text" value="12-01-2015"/>
End Date	<input type="text" value="12-31-2015"/>

Note that user1 does not provide the password for this user. This is to ensure that only the actual contractor can log in with this account

- b. Click Start to initiate the process.
 - c. Log out of the Self-Service UI.
3. Approve the workflow task as manager1:
- a. Log in to the Self-Service UI as `manager1`, with password `Welcome1`.

- b. Under My Group's Tasks, locate the Approve Contractor task and select Assign to Me.
- c. Approve Contractor is now listed under My Tasks.

Click the Details link.

- d. Check the form content. (It is the same content that you provided as `user1`, along with a Decision field.)

Select Accept and click Complete to finish the task.

- e. Log out of the Self-Service UI.

4. Verify that the contractor has been created in the XML file:

Open `openidm/samples/workflow/data/xmlConnectorData.xml` and note the addition of the new contractor entry. Note that there is no value for `<icf: __PASSWORD __/>`. Note that `user1` is the contractor's manager.

The following excerpt of the `xmlConnectorData.xml` shows a new user entry for bjensen, after the implicit synchronization to the XML file:

```
<ri: __ACCOUNT __>
...
  <ri:roles>openidm-authorized</ri:roles>
  <ri:firstname>Barbara</ri:firstname>
  <ri:manager>user1</ri:manager>
  <icf: __UID __>a02536cf-84b2-4d5c-a3ae-480f5e0899e9</icf: __UID __>
  <icf: __NAME __>bjensen</icf: __NAME __>
  <ri:email>bjensen@example.com</ri:email>
  <icf: __PASSWORD __/>
  <ri:lastname>Jensen</ri:lastname>
</ri: __ACCOUNT __>
```

- #### 5. Complete the password reset process:
- a. Check the inbox for the email address that you provided when you completed the initial form.

You should have received an email with the subject "Reset your password".

- b. Open the password reset email and click Password reset link.

The link takes you to the Self-Service UI, with the option to Reset Your Password.

- c. Enter a new password and confirmation password, submit the form.

The password that you enter here must comply with the default password policy for managed users, described in Section 17.1, "Enforcing Password Policy" in the *Integrator's Guide*.

- d. Click Return to Login Page and log in with the username that you provided when you completed the initial form, and the new password you have just set.

Notice the Welcome message under Notifications.

6. Verify that the password reset has been propagated to the XML file:

Open `openidm/samples/workflow/data/xmlConnectorData.xml` and note that the password for the contractor has been added to their entry.

If you declined the approval request, the user is not created in either data source.

12.2. Workflow Use Cases

This section describes a number of sample workflows, that demonstrate typical use cases for OpenIDM. The use cases, provided in `/path/to/openidm/samples/usecase`, work together to describe a complete business story, with the same set of sample data. Each of the use cases is integrated with the Self-Service UI.

These use cases use OrientDB as a repository by default. Alternative repository configuration files are provided in `/path/to/openidm/samples/usecase/db`. If you want to use one of these alternative repositories, remove the `repo.orientdb.json` file from the `conf/` directory of the use case you are testing (for example, `samples/usecase/usecase1/conf/repo.orientdb.json`) and copy the appropriate JDBC repository configuration files (`datasource.jdbc-default.json` and `repo.jdbc.json`) into that `conf/` directory. For more information on using an alternative repository, see Chapter 2, "Installing a Repository For Production" in the *Installation Guide*.

Each use case builds on the previous one. You *must* run the use cases in order, from use case 1 through 3, before you try the remaining use cases. Use cases 2 onwards depend on the `hr_data.ldif` file that you import and reconcile when you run use case 1.

The use cases assume an initial data set of twenty *ordinary* managed users in OpenIDM (user.0 - user.19). The users are divided as follows:

Users	Department	Manager	Employees	Contractors
user.0-user.4	Human Resources	user.0	user.0-user.3	user.4
user.5-user.9	Production Planning	user.5	user.5-user.8	user.9
user.10-user.14	Sales & Distribution	user.10	user.10-user.13	user.14
user.15-user.19	Treasury & Payments	user.15	user.15-user.18	user.19

In addition, the following *special* users are defined:

- `hradmin` - represents the human interaction of the HR department
- `systemadmin` - represents the human interaction of the populated systems (Business and Project)
- `superadmin` - represents the manager of the managers

Note

Note that the `curl` commands in this section use the secure port for OpenIDM (8443) and assume a self-signed certificate named `self-signed.crt`, located in the directory from which the command is launched. For instructions on using the self-signed certificate that is generated when OpenIDM first starts up, see Section 19.2.2, "Restrict REST Access to the HTTPS Port" in the *Integrator's Guide*.

12.2.1. Use Case 1 - Initial Reconciliation

This use case assumes an OpenDJ server and populates the managed user repository with users from OpenDJ.

To set up the sample, install and configure OpenDJ, as follows:

1. Download and install OpenDJ, as described in [To Install OpenDJ Directory Server With the GUI](#).

This sample assumes that OpenDJ is listening on port 1389, the standard LDAP port for users who cannot use privileged ports.

2. The use case assumes a user with DN `cn=Directory Manager` and password `password` who will bind to the directory server.
3. During the install, import the user data from the LDIF file `/path/to/openidm/samples/usecase/data/hr_data.ldif`.

The OpenDJ server now contains the users required for all the workflow use cases.

Procedure 12.1. Running Use Case 1

1. Start OpenIDM with the configuration for use case 1.

```
$ cd /path/to/openidm
$ ./startup.sh -p samples/usecase/usecase1
```

2. Run reconciliation to populate the managed user repository with the users from the OpenDJ server.

The validation rules in this workflow require a managed user's *manager* entry to exist before that user can be created. You must therefore run three consecutive reconciliation operations:

- The first reconciliation creates the `superadmin` user. This user has no manager entry, so is unaffected by the validation rules. Creation fails for the remaining users, because their manager does not yet exist.
- The second reconciliation creates the 12 users who have the `superadmin` user as their manager. Creation fails for the remaining users, because they require the 12 manager users to exist before they can be created.

- The third reconciliation creates the remaining 10 users, bringing the total number of users to 23.

The easiest way to run reconciliation is from the Admin UI:

1. Log in to the Admin UI (<https://localhost:8443/admin>) as the administrative user (`openidm-admin`) with password `openidm-admin`.
2. Select Configure > Mappings.
3. Click on the only configured mapping (`systemHRAccounts_managedUser`) and click Reconcile Now.

To run reconciliation from the command line, use the following command:

```
$ curl \
--cacert self-signed.crt \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request POST \
"https://localhost:8443/openidm/recon?_action=recon&mapping=systemHRAccounts_managedUser"
{
  "_id": "376b3290-24f0-47a9-8a9e-bba025536c39",
  "state": "ACTIVE"
}
```

3. Run reconciliation twice more to create all the users in the repository.
4. Query the managed users that were created by the three reconciliation operations.

In the Admin UI, select Manage > User and note the new entries in the User List.

Alternatively, run the following command to view the managed user entries over REST:

```
$ curl \
--cacert self-signed.crt \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"https://localhost:8443/openidm/managed/user?_queryId=query-all-ids"
{
  "result": [
    {
      "_id": "user.7",
      "_rev": "1"
    },
    {
      "_id": "user.3",
      "_rev": "1"
    },
    {
      "_id": "user.4",
      "_rev": "1"
    },
    ...
  ]
}
```

```
    "_id": "systemadmin",
    "_rev": "1"
  },
  {
    "_id": "hradmin",
    "_rev": "1"
  },
  {
    "_id": "superadmin",
    "_rev": "1"
  }
],
"resultCount": 23,
...
}
```

Your managed user repository should now contain 23 users. The default password of all the newly created users is `Passw0rd`.

5. Shut down OpenIDM before you proceed with the next use case.

```
$ cd /path/to/openidm
$ ./shutdown.sh
```

12.2.2. Use Case 2 - New User Onboarding

This use case demonstrates a new user onboarding process. The process can be initiated by any of the users created in the previous reconciliation process. In this example, we use `user.1` to initiate the process. `user.1` captures the details of a new user, and then submits the new user entry for approval by the prospective manager of that new user.

The use case includes three separate workflows - onboarding (creation of the new user), sunrise (new user start date) and sunset (user end date).

The use case also demonstrates email notification with the configuration of an external email service. You must configure the external email service, as described in Procedure 12.2, "Configuring Email Notification", *before you start the workflow*.

The use case demonstrates the OpenIDM Self-Service UI. If you deploy OpenIDM on the local system, you can access the UI at the following URL: <https://localhost:8443>.

Procedure 12.2. Configuring Email Notification

1. Start OpenIDM with the configuration for use case 2.

```
$ cd /path/to/openidm
$ ./startup.sh -p samples/usecase/usecase2
```

2. Configure the outbound email service.

Log in to the Admin UI (<https://localhost:8443/admin/>) as the default administrative user (`openidm-admin` with password `openidm-admin`).

Select Configure > System Preferences > Email and provide connection information for your preferred email host.

Gmail is configured by default but you must still supply credentials to use Gmail.

3. Log out of the Admin UI.
4. Change the notification email parameters in the workflow definition file. To edit the workflow definition file:
 1. Copy the workflow archive (`.bar`) file (`newUserCreate.bar`) to a temporary location, such as a `/tmp` directory:

```
$ cd /path/to/openidm/samples/usecase/usecase2/workflow
$ cp newUserCreate.bar /tmp/
```

2. Unzip the temporary workflow `.bar` file.

This step extracts the workflow definition file (`newUserCreate.bpmn20.xml`) and two xhtml templates required by the workflow:

```
$ unzip /tmp/newUserCreate.bar
Archive:  newUserCreate.bar
  inflating: nUCDecideApprovalForm.xhtml
  inflating: nUCStartForm.xhtml
  inflating: newUserCreate.bpmn20.xml
```

3. Edit the extracted workflow definition file (`newUserCreate.bpmn20.xml`). The email parameters are towards the end of this file:

```
$ cd /tmp
$ grep emailParams newUserCreate.bpmn20.xml
emailParams = [from : 'usecasetest@forgerock.com', to : 'notification@example.com'
,
...

```

Change the `from` and `to` parameters to reflect valid email addresses.

4. Zip up the amended workflow definition file, and the xhtml templates into a workflow `.bar` file.

```
$ zip newUserCreate.bar newUserCreate.bpmn20.xml nUCDecideApprovalForm.xhtml nUCStartForm.xhtml
updating: nUCDecideApprovalForm.xhtml (deflated 82%)
updating: nUCStartForm.xhtml (deflated 82%)
updating: newUserCreate.bpmn20.xml (deflated 85%)
```

5. Copy the new `.bar` file to the workflow directory, overwriting the existing `.bar` file.

```
$ cp /tmp/newUserCreate.bar /path/to/openidm/samples/usecase/usecase2/workflow
```

Procedure 12.3. Initiating the Onboarding Workflow

1. Start OpenIDM with the configuration for use case 2 (if you have not already done so when you configured the email service).

```
$ cd /path/to/openidm
$ ./startup.sh -p samples/usecase/usecase2
```

2. Log in to the Self-Service UI (<https://localhost:8443>) as `user.1` with password `Passw0rd`.

In this sample, any user who logs into the Self-Service UI can view the new User Onboarding Process workflow.

3. Click Details next to User Onboarding Process link and complete the fields for a sample new user.

Department. Specifies one of four departments that the new user will belong to (Human Resources, Production Planning, Sales & Distribution, or Treasury & Payments). The value you select here determines the *manager* of the new user, to which the request will be sent for approval. (See the previous table of users for a list of the managers of each department.)

User Type. Governs user access to specific accounts. If the User Type is Employee, the new user will have access to an account named Business. This access is represented as an attribute of the managed user entry in the OpenIDM repository, as follows: `accounts : ["Business"]`. If the User Type is Contractor, the new user will have no accounts associated with its managed user entry in OpenIDM.

Send Email Notification. Indicates whether an email should be sent to alert the manager of the new required approval. The email details used here are defined when you configure email notification, as described in Procedure 12.2, "Configuring Email Notification". If you select not to send an email notification, the notification is simply added to the OpenIDM repository, and appears when the manager logs into the UI.

4. Click Start to initiate the onboarding workflow.

This action sends the new user request to the corresponding *management* users (the department manager, as well as the `superadmin` user, who is an overall manager).

5. Log out of the UI, and log back in as the management user of the department that you selected when you completed the new user form. For example, if you selected Human Resources, log in as `user.0`, which simulates the management user for the HR department. All users have the password `Passw0rd`.


Notice that the management user now has an Onboarding Approval task in the queue of tasks assigned to that user's group.

Dashboard


MY TASKS

No tasks assigned.

MY GROUP'S TASKS

	Onboarding Approval Not Assigned	<input type="text" value="Assign to Me"/>	▶ Details
---	--	---	---------------------------

PROCESSES

	User onboarding process	▶ Details
---	--------------------------------	---------------------------

6. Select Assign to Me from list next to the Onboarding Approval task.

This action *claims* the task for `user.0`, removes it from the group queue, and places it in the list of pending tasks for `user.0`.

7. Select Details next to the Onboarding Approval task under My Tasks.

The complete new user request is displayed for the manager's approval. As the manager, you can add any information that was missing from the original request.

In addition, you can specify the following information for the new user.

- *Start Date*. Completing this field results in the user being created, with a `startDate` added to that user's managed user entry. The status of the user is `inactive`. This field is optional, and is used by the task scanner to trigger the Sunrise workflow.

- **End Date.** Completing this field results in the user being created, with an `endDate` added to that user's managed user entry. The field is optional, and is used by the task scanner to trigger the Sunset workflow.
- **Decision.** Selecting Reject here terminates the workflow and sends a notification to the user who initiated the workflow. Selecting Accept creates the managed user entry in OpenIDM. The password of the new user is `Passw0rd`.

Two notifications are created when the request is accepted - one for the user who initiated the workflow, and one for the newly created user. The notifications are visible in the UI after login. If you selected email notification, one email is sent to the user that you defined when you configured email notification, as described in Procedure 12.2, "Configuring Email Notification".

8. At the bottom of the form, there is an option either to Requeue the request or to Complete it. Click Complete.

If you click Requeue here, the task is removed from the list of that user's tasks, and returned to the list of tasks pending for that group. The task can then be claimed by any member of that group.

When the new user request has been approved, the user is created in the OpenIDM repository. If you did not include a Start Date in the manager approval, you should now be able to log into the UI with the details of the new user. If you included a Start Date, you need to complete the sunrise workflow before the user account is active (which will enable you to log in as this user).

Procedure 12.4. Initiating the Sunrise Workflow

If a sunrise date is specified for the new user, the user is created in the repository, with an `inactive` account status.

- To trigger the sunrise workflow (which activates the account), enable the sunrise task scanning schedule. The schedule is disabled by default.

Modify the schedule configuration file (`schedule-taskscan_sunrise.json`), setting the `enabled` property to `true`.

```
$ cd /path/to/openidm
$ grep "enabled" samples/usecase/usecase2/conf/schedule-taskscan_sunrise.json
"enabled" : true,
```

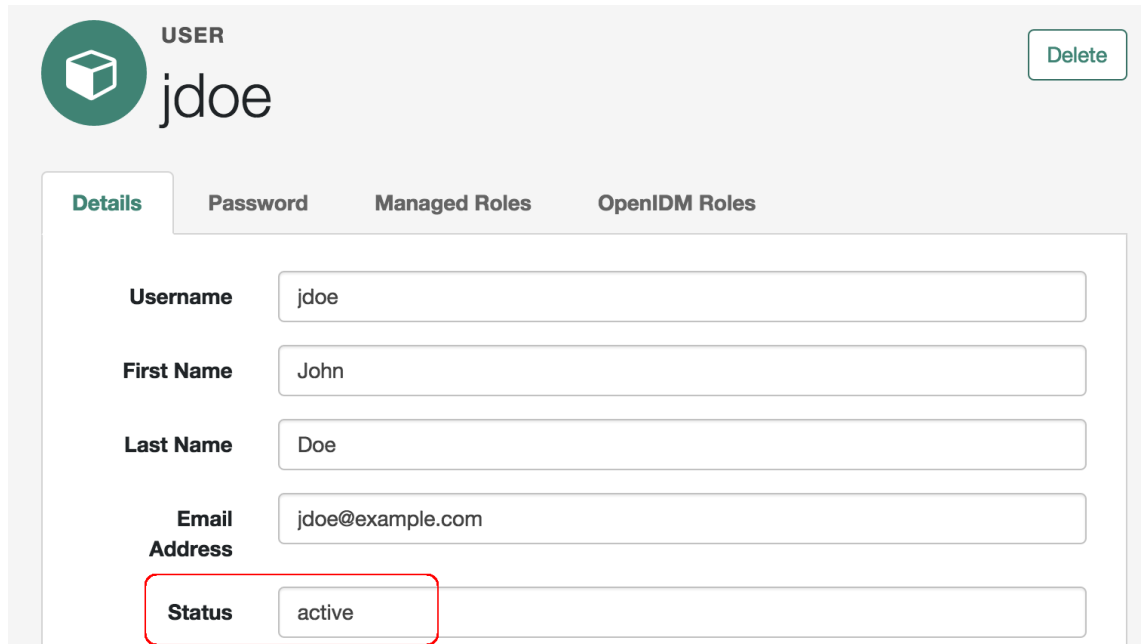
The scan runs every minute, and checks the repository for users that have a sunrise date that is anything up to one day after the current date. When the scan is triggered, it locates the newly created user and starts the sunrise workflow on this user. The workflow takes the following actions:

- Changes the account status of the user to `active`.

You can check that this part of the workflow has completed by looking at the user's account in the Admin UI:

1. Log in to the Admin UI as `openidm-admin` with password `openidm-admin`, then select Manage > User.
2. Click on the account of the new user that was created in the previous section and note their Status.

The following image shows the status for the new user `jdoue`, who was created by the previous workflow:

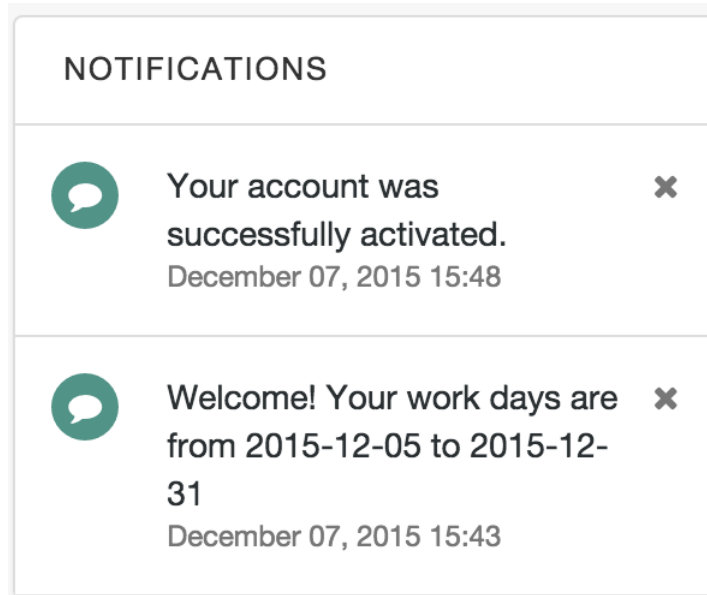


The screenshot displays the user management interface for a user named 'jdoue'. At the top left, there is a green circular profile icon with a white cube inside, followed by the text 'USER' and 'jdoue'. A 'Delete' button is located in the top right corner. Below the user information, there are four tabs: 'Details', 'Password', 'Managed Roles', and 'OpenIDM Roles'. The 'Details' tab is selected and shows the following fields:

Username	jdoue
First Name	John
Last Name	Doe
Email Address	jdoue@example.com
Status	active

The 'Status' field, which contains the value 'active', is highlighted with a red rectangular border.

- Generates a notification for the new user, which is visible when the user logs into the Self-Service UI.



Procedure 12.5. Initiating the Sunset Workflow

If a sunset date is set for the new user, you can trigger the sunset workflow to deactivate the user account when the end of his work period is reached.

1. To trigger the sunset workflow, enable the sunset task scanning schedule. The schedule is disabled by default.

Modify the schedule configuration file (`schedule-taskscan_sunset.json`), setting the `enabled` property to `true`.

```
$ cd /path/to/openidm
$ grep "enabled" samples/usecase/usecase2/conf/schedule-taskscan_sunset.json

"enabled" : true,
```

The scan runs every minute, and checks the repository for users that have a sunset date that is anything up to one day after the current date. When the scan is triggered, it locates users whose contracts are about to end, and starts the sunset workflow on these users. When the workflow is initiated, it assigns a task to the manager of the affected user. In our example, the task is assigned to `user.0`.

2. When the sunset schedule has been enabled, log in to the Self-Service UI as `user.0` (with password `Passw0rd`). If the user's sunset date is within one day of the current date, a Contract Termination task becomes available under the manager's My Group's Tasks section.

Select the contract termination task and click Details.

3. In the Decision field, select either Accept termination or Modify date, then click Complete.

When you accept the termination, the user's account status is set to **inactive** and the HR administrative user receives a notification to that effect, the next time that user logs into the UI. The deactivated user is no longer able to log into the UI.

If you select to modify the date, the sunset date of that user is changed to the value that you specify in the End Date field on that form. The management user receives a UI notification that the employee's contract has been extended.

4. Shut down OpenIDM before you proceed with the next use case.

```
$ cd /path/to/openidm
$ ./shutdown.sh
```

12.2.3. Use Case 3 - User Access Request

This use case simulates a user access request, with two levels of approval for the request.

If you want to use email notification with this workflow, start OpenIDM with the configuration for Use Case 3, then follow the instructions in Procedure 12.2, "Configuring Email Notification", substituting and `usecase3/workflow/accessRequest.bpmn20.xml` for the file described in that procedure.

1. Start OpenIDM with the configuration for use case 3, if you have not already started the server to configure email notification.

```
$ cd /path/to/openidm
$ ./startup.sh -p samples/usecase/usecase3
```

2. Log into the Self-Service UI as **user.1** with password **Passw0rd**.

user.1 belongs to the HR department and, in this workflow, is requesting access to a Project system.

3. Click Details next to the Access request process and click Start to start the workflow.

A User Access Request appears in the list of tasks for **user.1**.

4. Click Details next to the User Access Request task.

The resulting form indicates the various systems to which the user may request access.

Access to Business system. This field reflects the current value of the **accounts** property for that user in the repository. If the value includes **Business** this field is True.

Access to Project system. Set this field to True to request Project access for **user.1**.

Send Email Notification. Set to True to send an email to alert the manager of the new access request. The email details used here are defined when you configure email notification, as described in Procedure 12.2, "Configuring Email Notification". If you select not to send an email notification, the notification appears when the manager logs into the UI.

Select either Cancel, to terminate the process, or Request, to start a user task, assigned to the manager of the user who is requesting access (`user.0` in this example), and select Complete.

5. Log out of the Self-Service UI and log back in as the manager (`user.0` with password `Passw0rd`).
6. Under My Group's Tasks, locate the User Access Request Approval task and select Assign to me.

Note that the User Access Request Approval task has moved under My Tasks.


7. Click Details next to the User Access Request Approval task.
8. The details of the access request are displayed. The manager is able to modify the access rights. Select Accept or Reject to approve or deny the request.

Rejecting the request results in a notification being sent to the user who made the request. If you have enabled email notification, a single email is sent to the account that you specified when you configured email notification.

Accepting the request initiates a second approval task, assigned to the `systemadmin` user.

Dashboard

MY TASKS

 **User Access Request Approval** ▼ Details
4 minutes

Username

First Name

Last Name

Email

Access to Business system:

Access to Project system:

Decision

[Close](#) [Requeue](#) [Complete](#)

Click Complete to complete the task.

9. Log out of the UI and log in as the `systemadmin` user (with password `Passw0rd`).

This user now has one User Access Request Approval task in his queue.

10. Select the task and click Details.

This task interface is similar to that of the task that was assigned to the manager.

Rejecting the request results in a notification being sent to the user who made the request.

Accepting the request updates the managed/user record in OpenIDM, to reflect the approved access changes.

If you have enabled email notification, a single email is sent to the account defined when you configured the external email service (Procedure 12.2, "Configuring Email Notification"), indicating whether the request has been accepted or rejected.

Note that this sample includes an *escalation* step that is attached to the manager approval task. If the manager does not complete assessment of the user task within ten minutes of its initiation, a new user task is created and assigned to the `superadmin` user. This task has the same interface and functionality as the task assigned to the manager. Accordingly, when the `superadmin` user completes the task, the execution is passed to the `systemadmin` user for approval.

Shut down OpenIDM before you proceed with the next use case.

```
$ cd /path/to/openidm
$ ./shutdown.sh
```

12.2.4. Use Case 4 - Orphan Account Detection

This use case demonstrates two asynchronous tasks, started from a reconciliation process:

- Detecting orphan accounts on a target object set
- Handling ambiguous results during correlation

This use case relies on a customized synchronization configuration (mapping) file, named `syncManagedBusiness.json`, in the `/path/to/openidm/samples/usecase/usecase4/conf` directory.

This file defines a mapping (`recon_managedUser_systemBusiness`) between a source (managed users) and a target object set. The target object set is defined in the file `samples/usecase/usecase4/data/business.csv`. The `business.csv` file includes all users from the initial reconciliation (described in Section 12.2.1, "Use Case 1 - Initial Reconciliation"). These users are categorized as `employees`, and therefore include the property `"accounts" : ["Business"]` in their managed user entry (see Section 12.2.2, "Use Case 2 - New User Onboarding" for an explanation of the User Type).

The mapping includes the following `"validSource"` field:

```
"validSource" : {
  "type" : "text/javascript",
  "file" : "script/isSourceValidBusiness.js"
},
```

This field references a script which specifies that only those users who are employees are taken into account during the reconciliation.

In addition, the `business.csv` file includes the following users:

- `user.50`. This user is defined *only* in the `.csv` file, and not in the managed/user repository. When a reconciliation operation is run, this user is detected as an *orphan account*. The orphan account workflow is triggered when an "UNQUALIFIED" or "UNASSIGNED" situation is encountered, as indicated in this section of the mapping:

```
{
  "situation" : "UNQUALIFIED",
  "action" : {
    "workflowName" : "orphanAccountReport",
    "type" : "text/javascript",
    "file" : "workflow/triggerWorkflowFromSync.js"
  }
},
{
  "situation" : "UNASSIGNED",
  "action" : {
    "workflowName" : "orphanAccountReport",
    "type" : "text/javascript",
    "file" : "workflow/triggerWorkflowFromSync.js"
  }
}
}
```

- **user.33.** This user has a `userName` attribute of `user.3` (which is the same as the `userName` attribute of the user, `user.3`). The correlation query of the reconciliation operation is based on the `userName` attribute. During the correlation query, two candidate users are therefore correlated with the same managed user (`user.3`), and the result is ambiguous. The manual match workflow is triggered when an `AMBIGUOUS` situation is encountered, as indicated in this section of the mapping:

```
{
  "situation" : "AMBIGUOUS",
  "action" : {
    "workflowName" : "manualMatch",
    "type" : "text/javascript",
    "file" : "workflow/triggerWorkflowFromSync.js"
  }
}
}
```

1. Before you start with this use case, rename the mapping file to `sync.json`.

```
$ cd /path/to/openidm/samples/usecase/usecase4/conf
$ mv syncManagedBusiness.json sync.json
```

2. Start OpenIDM with the configuration for use case 4.

```
$ cd /path/to/openidm
$ ./startup.sh -p samples/usecase/usecase4
```

3. Reconcile the managed user repository with the CSV file, either by using the Admin UI or over the command line.

To use the Admin UI, log in to the Admin UI (<https://localhost:8443/admin/>) as `openidm-admin` with password `openidm-admin`. Select Configure > Mappings, click on the mapping `recon_managedUser_systemBusiness` and click Reconcile Now.

To use the command line, run the following command:

```
$ curl \
--cacert self-signed.crt \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request POST \
"https://localhost:8443/openidm/recon?action=recon&mapping=recon_managedUser_systemBusiness"
```

When the reconciliation operation finds the ambiguous entry (`user.3`) and the orphan entry (`user.50`) in the CSV file, two asynchronous workflows are launched (`manualMatch` and `orphanAccountReport`), as indicated in the mapping file, described previously.

4. Log in to the Self-Service UI (<https://localhost:8443>) as the `systemadmin` user, with password `Passw0rd`.
5. Next to the Manual Linking Task in the My Tasks list, click Details.

The *Possible targets* field presents a list of target entries to which the ambiguous record can be linked. In this example, `user.3 - Atrc, Aaron` and `user.33 - Atrc, Aaron` are the two candidate users found in the target object set by the correlation query. When you select one of these values, the workflow manually links the managed user (`user.3`) to the selected user.

Click complete to finish the manual account linking task.

If you select Ignore, here, no action is taken (no link is created), and the workflow terminates.

6. Next to the Orphan Account Task in the My Tasks list, click Details.

The *Link to* field enables you to enter an existing managed user ID to which this orphan account should be linked. For the purposes of this example, enter `user.5`.

In the Decision field, select Link to link the orphan account to the ID that you entered in the previous step. Click Complete to complete the task.

Selecting Delete here deletes the user from the target object set (the CSV file in this case) and terminates the workflow.

7. Shut down OpenIDM before you proceed with the next use case.

```
$ cd /path/to/openidm
$ ./shutdown.sh
```

Note

Use Case 5 has been removed from the sample use cases.

12.2.5. Use Case 6 - Password Change Reminder

This use case demonstrates using the task scanner to trigger a password change reminder workflow for managed users.

In this example, each managed user entry in OpenIDM has a dedicated attribute, `lastPasswordSet`, that stores the date on which the password was last changed. The value of this attribute is updated by an `onStore` script, defined in the managed user configuration file (`conf/managed.json`), as follows:

```
"onStore" : {
  "type" : "text/javascript",
  "file" : "script/onStoreManagedUser.js"
},
```

When a new password is stored for a user, the script sets the date on which this change was made. The task scanner periodically scans the `lastPasswordSet` attribute, and starts the workflow if the password was changed more than an hour ago. This condition is configured in the schedule configuration file (`schedule-taskscan_passwordchange.json`):

```
$ cd /path/to/openidm
$ more samples/usecase/usecase6/conf/schedule-taskscan_passwordchange.json

...
"condition" : {
  "before" : "${Time.now - 1h}"
}
,
....
```

Obviously, in a real deployment, the period between required password changes would be longer, and this value would need to be set accordingly. For the purposes of testing this use case, you might want to set the value to a shorter period, such as `"${Time.now - 1m}"`, which will send the notification one minute after a password change.

By default, the workflow sends notifications to the user entry, visible when the user logs into the UI. If you want notifications sent by email, configure the external email service, as follows:

1. Set up outbound email, as described in Procedure 12.2, "Configuring Email Notification".
2. Enable email notification in the script file that starts the workflow (`samples/usecase/usecase6/script/passwordchange.js`). For example:

```
$ cd /path/to/openidm
$ more samples/usecase/usecase6/script/passwordchange.js

/*global objectID*/

(function () {
  var params = {
    "userId" : objectID,
    "emailEnabled" : "true",
    "_key": "passwordChangeReminder"
  };
```

3. Make sure that all managed users have a valid email address as the value of their `mail` attribute.

The task scanning schedule is disabled by default. To test this use case, follow these steps:

1. Enable the task scanning schedule by setting `enabled` to `true` in the schedule configuration file (`schedule-taskscan_passwordchange.json`).

```
$ cd /path/to/openidm
$ more samples/usecase/usecase6/conf/schedule-taskscan_passwordchange.json
{
  "enabled" : true
,
...
}
```

2. Start OpenIDM with the configuration for use case 6.

```
$ cd /path/to/openidm
$ ./startup.sh -p samples/usecase/usecase6
```

3. Log in to the Self-Service UI as any of the users listed in the introduction to this section (for example, `user.4`, with password `Passw0rd`).

You will see a password expiration notice.

If you ignore the change for five minutes, you will see a second warning, that the password is about to expire.

If you ignore that second warning for another two minutes, the user's account is deactivated.

4. To avoid the second notification, or the account deactivation, you can change the user password through the UI, as follows:
 - a. Log in to the UI as the user whose password you want to change and click Change Password at the top right of the page.
 - b. Enter a new password that conforms to the requirements of the password policy.
 - c. Enter the old password (in this case `Passw0rd`).

Chapter 13

Google Sample - Connecting to Google With the Google Apps Connector

The Google Apps Connector that enables you to interact with Google's web applications.

This sample demonstrates the creation of users and groups on an external Google system, using OpenIDM's REST interface. The sample requires that you have a Google Apps account. Obtaining a Google Apps account is described in the [Google documentation](#).

13.1. Before You Start

To set up OpenIDM to connect to your Google Apps account, you must have a Google Apps project (or create a new project) that authorizes consent for OpenIDM.

1. Log in to the Google Apps Developers Console (at <https://console.developers.google.com/start>) and update your project or create a new project for OpenIDM.
2. Enable the following APIs for your OpenIDM project:
 - Admin SDK API
 - Enterprise License Manager API
3. Set up an OAuth2 Client ID.

The Google Apps connector uses OAuth2 to authorize the connection to the Google service. Set up an OAuth2 Client ID as follows:

- a. In the Google Apps Developers Console, select Credentials > New Credentials > OAuth Client ID.
- b. Click Configure Consent Screen and enter a Product Name.

This is the name that will be shown for all applications registered in this project.

For the purposes of this example, we use the Product Name **OpenIDM**.

Click Save.

- c. Select Credentials > OAuth Client ID > Web application.

Under Authorized redirect URIs, enter the callback URL (the URL at which your clients will access your application). The default OpenIDM callback URL is <https://localhost:8443/admin/oauth.html>. Click Create to set up the callback URL.

Click Create again to set up the client ID.

This step generates an OAuth Client ID and Client, similar to the following:

OAuth client

Here is your client ID

```
480173612956-2cqavtu0nfde0m0cpa6qq3kt4b71im0l.apps.googleusercontent.com
```

Here is your client secret

```
3CIHYFrN03npt-PTqkBh1N2j
```

Copy and paste these values into a text file as you will need them when you configure the Google Apps connector.

13.2. Configuring the Google Apps Connector

This procedure uses the OpenIDM Admin UI to set up the Google Apps connector.

1. To configure the connector, start OpenIDM with the Google Apps sample configuration:

```
$ cd /path/to/openidm
$ ./startup.sh -p samples/google-connector
Executing ./startup.sh...
Using OPENIDM_HOME: /path/to/openidm
Using PROJECT_HOME: /path/to/openidm/samples/google-connector/
Using OPENIDM_OPTS: -Xmx1024m -Xms1024m
Using LOGGING_CONFIG: -Djava.util.logging.config.file=/path/to/openidm/conf/logging.properties
Using boot properties at /path/to/openidm/samples/google-connector/conf/boot/boot
.properties
-> OpenIDM ready
```









2. Log in to the Admin UI at the URL <https://localhost:8443/admin> as the default administrative user (`openidm-admin`) with password `openidm-admin`.

This URL reflects the host on which OpenIDM is installed and corresponds to the callback URL that you specified in the previous section. The URL must be included in the list of Authorized redirect URIs for your project.

3. Select Configure > Connectors and click on the Google Apps connector.
4. On the Details tab, set the Enabled field to True.
5. Enter the Oauth2 Client ID and Client Secret that you obtained in the previous section.
6. Click Save Connector Changes.
7. You are redirected to Google's Login page.

When you have logged in, Google requests that you allow access from your project, in this case, OpenIDM.

▼ OpenIDM would like to:

-
- | | | |
|---|---|---|
|  | View and manage the provisioning of users on your domain |  |
|  | View and manage the provisioning of groups on your domain |  |
|  | View and manage Google Apps licenses for your domain |  |
|  | View and manage organization units on your domain |  |
-

By clicking Allow, you allow this app and Google to use your information in accordance with their respective terms of service and privacy policies. You can change this and other [Account Permissions](#) at any time.

Click Allow.

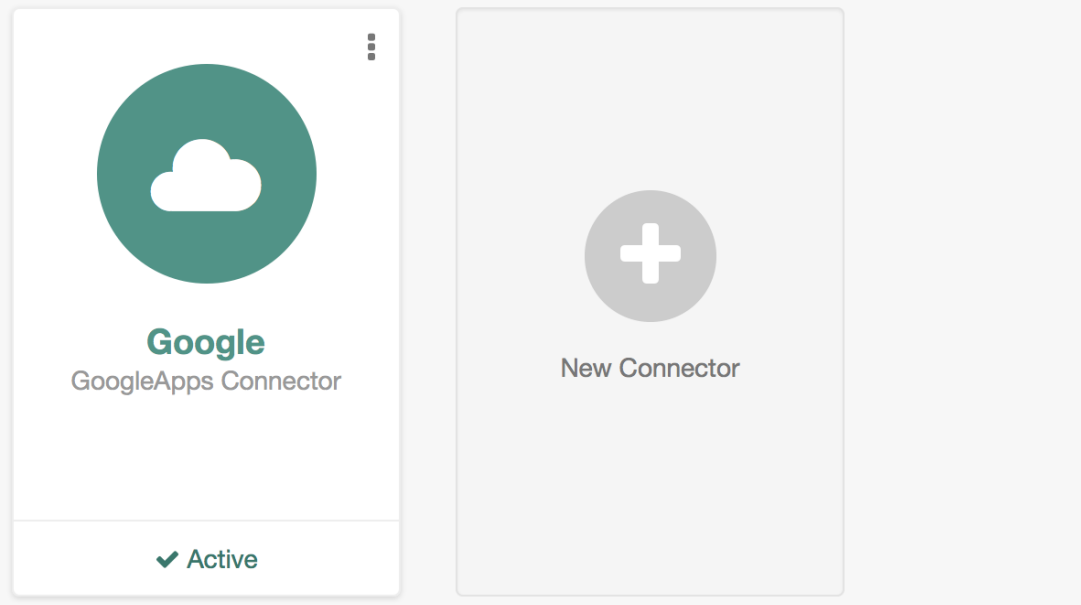
If you click Deny here, you will need to return to the Connector Configuration > Details tab in the Admin UI and save your changes again.

When you allow access, you are redirected to the Connectors page in the OpenIDM Admin UI, where the Google Apps Connector should now be Active.

Connectors

Connectors provide interfaces to external systems, databases, directory services, and more.

+ New Connector



13.3. Running the Google Apps Sample

This procedure uses create, read, update, and delete (CRUD) operations to the Google resource, to verify that the connector is working as expected. The procedure uses a combination of REST commands, to manage objects on the Google system, and the Admin UI, to manage reconciliation from the Google system to the managed user repository.

The sample configuration has one mapping *from* the Google system *to* the managed user repository.

All of the commands shown here assume that your domain is `example.com`. Adjust the examples to manage your domain.

1. Create a user entry on your Google resource, over REST.

When you create resources for Google, note that the equals (=) character cannot be used in any attribute value.

The following command creates an entry for user **Sam Carter**:

```
$ curl \
--cacert self-signed.crt \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--header "Content-Type: application/json" \
--request POST \
--data '{
  "__NAME__": "samcarter@example.com",
  "__PASSWORD__": "password",
  "givenName": "Sam",
  "familyName": "Carter",
  "agreedToTerms": true,
  "changePasswordAtNextLogin": false
}' \
"https://localhost:8443/openidm/system/google/__ACCOUNT__?_action=create"
{
  "_id": "103567435255251233551",
  "_rev": "\iwpzoDgSq9BJw-XzORg0bILYPVc/LWHPMXXG8M0cjQAPITM95Y636cM\"",
  "orgUnitPath": "/",
  "isAdmin": false,
  "fullName": "Sam Carter",
  "customerId": "C02rsqddz",
  "relations": null,
  "nonEditableAliases": null,
  "suspensionReason": null,
  "includeInGlobalAddressList": true,
  "givenName": "Sam",
  "addresses": null,
  "isDelegatedAdmin": false,
  "changePasswordAtNextLogin": false,
  "isMailboxSetup": true,
  "__NAME__": "samcarter@example.com",
  "agreedToTerms": true,
  "externalIds": null,
  "ipWhitelisted": false,
  "aliases": null,
  "lastLoginTime": [
    "1970-01-01T00:00:00.000Z"
  ],
  "organizations": null,
  "suspended": false,
  "deletionTime": null,
  "familyName": "Carter",
  "ims": null,
  "creationTime": [
    "2016-02-02T12:52:30.000Z"
  ],
  "thumbnailPhotoUrl": null,
  "emails": [
    {
      "address": "samcarter@example.com",
```

```
    "primary": true
  }
},
"phones": null
}
```

Note the ID of the new user (`103567435255251233551` in this example). You will need this ID for the update commands in this section.

2. Reconcile the Google resource with the managed user repository.

This step should create the new user, Sam Carter (and any other users in your Google resource) in the OpenIDM managed user repository.

To run reconciliation follow these steps:

- a. In the Admin UI, select Configure > Mappings.
 - b. Click on the sourceGoogle__ACCOUNT__managedUser mapping, and click Reconcile Now.
 - c. Select Manage > User and verify that the user Sam Carter has been created in the repository.
3. Update Sam Carter's phone number in your Google resource by sending a PUT request with the updated data, and specifying the user `_id` in the request:

```
$ curl \
--cacert self-signed.crt \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--header "Content-Type: application/json" \
--request PUT \
--header "If-Match : *" \
--data '{
  "_NAME_": "samcarter@example.com",
  "_PASSWORD_" : "password",
  "givenName": "Sam",
  "familyName": "Carter",
  "agreedToTerms": true,
  "changePasswordAtNextLogin" : false,
  "phones" :
  [
    {
      "value": "1234567890",
      "type": "home"
    },
    {
      "value": "0987654321",
      "type": "work"
    }
  ]
}' \
"https://localhost:8443/openidm/system/google/__ACCOUNT__/103567435255251233551"
{
  "_id": "103567435255251233551",
```

```
"_rev": "\"iwpzoDgSq9BJw-Xz0Rg0bILYPVc/vfSjGht-STUuto4lM_4ES09izR4\""  
,  
...  
"emails": [  
  {  
    "address": "samcarter@example.com",  
    "primary": true  
  }  
],  
"phones": [  
  {  
    "value": "1234567890",  
    "type": "home"  
  },  
  {  
    "value": "0987654321",  
    "type": "work"  
  }  
]  
}
```

4. Read Sam Carter's entry from your Google resource by including his `_id` in the URL:

```
$ curl \br/>--cacert self-signed.crt \  
--header "X-OpenIDM-Username: openidm-admin" \  
--header "X-OpenIDM-Password: openidm-admin" \  
--request GET \  
"https://localhost:8443/openidm/system/google/_ACCOUNT_/103567435255251233551"  
{  
  "_id": "103567435255251233551",  
  "_NAME_": "samcarter@example.com"  
,  
...  
  "phones": [  
    {  
      "value": "1234567890",  
      "type": "home"  
    },  
    {  
      "value": "0987654321",  
      "type": "work"  
    }  
  ]  
}
```

5. Create a group entry on your Google resource:


```
$ curl \
--cacert self-signed.crt \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--header "Content-Type: application/json" \
--request POST \
--data '{
  "__NAME__": "testGroup@example.com",
  "__DESCRIPTION__": "Group used for google-connector sample.",
  "name": "TestGroup"
}' \
"https://localhost:8443/openidm/system/google/__GROUP__?_action=create"

{
  "_id": "00meukdy40gpg98",
  "_rev": "\"iwpzoDg5q9BJw-XzORg0bILYPVc/LLhHx2p1MJPKeY1-h6eX_OVDi4c\"",
  "adminCreated": true,
  "__NAME__": "testgroup@example.com",
  "aliases": null,
  "nonEditableAliases": null,
  "__DESCRIPTION__": "Group used for google-connector sample.",
  "name": "TestGroup",
  "directMembersCount": 0
}
```

6. Add Sam Carter to the test group you have just created. Include the **Member** endpoint, and Sam Carter's **_id** in the URL. Specify the **_id** of the group you created as the value of the **groupKey** in the JSON payload:

```
$ curl \
--cacert self-signed.crt \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--header "Content-Type: application/json" \
--request PUT \
--data '{
  "groupKey": "00meukdy40gpg98",
  "role": "MEMBER",
  "__NAME__": "samcarter@example.com",
  "email": "samcarter@example.com",
  "type": "MEMBER"
}' \
"https://localhost:8443/openidm/system/google/Member/103567435255251233551"

{
  "_id": "00meukdy40gpg98/samcarter@example.com",
  "_rev": "\"iwpzoDg5q9BJw-XzORg0bILYPVc/CPNpkRnowkGWRvNqVUK9ev6gQ90\"",
  "__NAME__": "00meukdy40gpg98/samcarter@example.com",
  "role": "MEMBER",
  "email": "samcarter@example.com",
  "type": "USER",
  "groupKey": "103567435255251233551"
}
```

7. Read the group entry by specifying the group **_id** in the request URL. Notice that the group has one member (**"directMembersCount": 1**).

```
$ curl \
--cacert self-signed.crt \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"https://localhost:8443/openidm/system/google/__GROUP__/00meukdy40gpg98"

{
  "_id": "00meukdy40gpg98",
  "_rev": "\"iwpzoDgSq9BJw-XzORg0bILYPVc/chUdq5m5_cycV2G4sd17ZKAF75A\"",
  "adminCreated": true,
  "__NAME__": "testgroup@example.com",
  "aliases": null,
  "nonEditableAliases": [
    "testGroup@example.test-google-a.com"
  ],
  "__DESCRIPTION__": "Group used for google-connector sample.",
  "name": "TestGroup",
  "directMembersCount": 1
}
```

8. Delete the group entry.

```
$ curl \
--cacert self-signed.crt \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request DELETE \
"https://localhost:8443/openidm/system/google/__GROUP__/00meukdy40gpg98"

{
  "_id": "00meukdy40gpg98",
  "_rev": "\"iwpzoDgSq9BJw-XzORg0bILYPVc/chUdq5m5_cycV2G4sd17ZKAF75A\"",
  "adminCreated": true,
  "__NAME__": "testgroup@example.com",
  "aliases": null,
  "nonEditableAliases": [
    "testGroup@example.com.test-google-a.com"
  ],
  "__DESCRIPTION__": "Group used for google-connector sample.",
  "name": "TestGroup",
  "directMembersCount": 1
}
```

The delete request returns the complete group object.

9. Delete Sam Carter, to return your Google resource to its original state.

```
$ curl \
--cacert self-signed.crt \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request DELETE \
"https://localhost:8443/openidm/system/google/__ACCOUNT__/103567435255251233551"

{
  "_id": "103567435255251233551",
  "_rev": "\"iwpzoDgSq9BJw-XzORg0bILYPVc/ah6xBLujMAHieSwSisPa1CV6T3Q\"",
  "orgUnitPath": "/",
  "isAdmin": false,
}
```

```
"fullName": "Sam Carter",
"customerId": "C02rsqddz",
"relations": null,
"nonEditableAliases": [
  "samcarter@example.com.test-google-a.com"
],
"suspensionReason": null,
"includeInGlobalAddressList": true,
"givenName": "Sam",
"addresses": null,
"isDelegatedAdmin": false,
"changePasswordAtNextLogin": false,
"isMailboxSetup": true,
"__NAME__": "samcarter@example.com",
"agreedToTerms": true,
"externalIds": null,
"ipWhitelisted": false,
"aliases": null,
"lastLoginTime": [
  "1970-01-01T00:00:00.000Z"
],
"organizations": null,
"suspended": false,
"deletionTime": null,
"familyName": "Carter",
"ims": null,
"creationTime": [
  "2016-02-02T12:52:30.000Z"
],
"thumbnailPhotoUrl": null,
"emails": [
  {
    "address": "samcarter@example.com",
    "primary": true
  }
],
"phones": [
  {
    "value": "1234567890",
    "type": "home"
  },
  {
    "value": "0987654321",
    "type": "work"
  }
]
}
```

In this sample, you used the Google Apps connector to add and delete user and group objects in your Google application, and to reconcile users from your Google application to the OpenIDM managed user repository. You can expand on this sample by customizing the connector configuration to provide additional synchronization functionality between OpenIDM and your Google applications. For more information on configuring connectors, see Chapter 13, *"Connecting to External Resources"* in the *Integrator's Guide*.

Chapter 14

Connecting to Salesforce With the Salesforce Connector

The Salesforce connector enables provisioning, reconciliation, and synchronization between Salesforce and the OpenIDM repository.

This sample demonstrates the creation and update of users from OpenIDM to Salesforce, and from Salesforce to OpenIDM. You can use either the Admin UI, or the command line to run this sample. Both methods are outlined in the sections that follow.

14.1. Before you Start

This sample requires that you have a Salesforce account, and a Connected App with OAuth enabled. For more information about Connected Apps, see the [Connected Apps Overview](#) in the Salesforce documentation.

To set up a Connected App for OpenIDM, follow these steps:

1. Log in to salesforce.com with your Salesforce credentials.
2. Click *Setup* in the top right corner.
3. In the left hand menu, under *Build*, expand the *Create* item and click *Apps*.
4. On the right hand panel, scroll down to *Connected Apps* and click *New*.
5. In the *New Connected App* panel, enter the following Basic Information:
 - *Connected App Name*. Enter a name that you will recognize as the OpenIDM App, for example, **OpenIDM**.
 - *API Name*. This field defaults to the Connected App Name, but you can change it. Note that the Application API Name can only contain underscores and alphanumeric characters. The name must be unique, begin with a letter, not include spaces, not end with an underscore, and not contain two consecutive underscores.
 - *Contact Email*. Enter the email address of the person responsible for this Connected App within your organization.
6. Select *Enable OAuth Settings* and enter the following information:

- **Callback URL.** Enter the OpenIDM URL, to which the requested OAuth token will be sent, for example <https://localhost:8443/admin/oauth.html>.
- **Selected OAuth Scopes.** Click the *Add* button to add the following *Available Auth Scopes* to the *Selected OAuth Scopes* column:
 - Access and manage your data
 - Access your basic information
 - Perform requests on your behalf at any time

New Connected App

[Help for this Page](#)

To publish an app, you need to be using a Developer Edition organization with a namespace prefix chosen.

Basic Information

|= Required Inform

Connected App Name

API Name

Contact Email

Contact Phone

Logo Image URL
Upload logo image or [Choose one of our sample logos](#)

Icon URL
[Choose one of our sample logos](#)

Info URL

Description

API (Enable OAuth Settings)

Enable OAuth Settings

Callback URL

Use digital signatures

Selected OAuth Scopes	Available OAuth Scopes	Add	Selected OAuth Scopes
<ul style="list-style-type: none"> Access and manage your Chatter data (chatter_api) Access custom permissions (custom_permissions) Allow access to your unique identifier (openid) Full access (full) Provide access to custom applications (visualforce) Provide access to your data via the Web (web) 	<ul style="list-style-type: none"> Access and manage your data (api) Access your basic information (id, profile, email, address, phone) Perform requests on your behalf at any time (refresh_token, offline_access) 	<input type="button" value="Add"/> <input type="button" value="Remove"/>	<ul style="list-style-type: none"> Access and manage your data (api) Access your basic information (id, profile, email, address, phone) Perform requests on your behalf at any time (refresh_token, offline_access)

You can leave the remaining fields blank.

- Click *Save* to create the new Connected App.
- The next window displays your new Connected App.

Under the *API (Enable OAuth Settings)* item, the Consumer Key and a link to the Consumer Secret are displayed.

Click the link to reveal the Consumer Secret.

▼ API (Enable OAuth Settings)			
Consumer Key	3MVG98dostKihXN7Is8Q0g5q1	PdB5f5ATwmaMuWxl	Consumer Secret 485 425
Selected OAuth Scopes	Access your basic information (id, profile, email, address, phone) Access and manage your data (api) Perform requests on your behalf at any time (refresh_token, offline_access)		Callback URL https://localhost:8443/admin/oauth.html

OpenIDM requires the Consumer Key and Secret to obtain an access token and a refresh token for access to salesforce.com.

Copy and paste both the key and the secret into a file for use when you set up the sample.

- To demonstrate the reconciliation of users from Salesforce to OpenIDM, your Salesforce organization should contain at least a few users. Add these users now if your Salesforce organization is empty.

14.2. Install the Sample

Prepare OpenIDM as described in Section 1.3, "Preparing the Server", then start OpenIDM with the configuration for the Salesforce sample.

```
$ cd /path/to/openidm
$ ./startup.sh -p samples/salesforce-connector
```

14.3. Running the Sample by Using the Admin UI

The Admin UI is the recommended way to test this sample.

- Log in to the Admin UI at the URL <https://localhost:8443/admin> as the default administrative user (`openidm-admin`) with password `openidm-admin`.

Warning

To protect your deployment in production, change the default administrative password. To do so, navigate to the Self-Service UI at <https://localhost:8443/> and click Change Password.

The Resources tab shows the Salesforce connector, which is currently disabled.

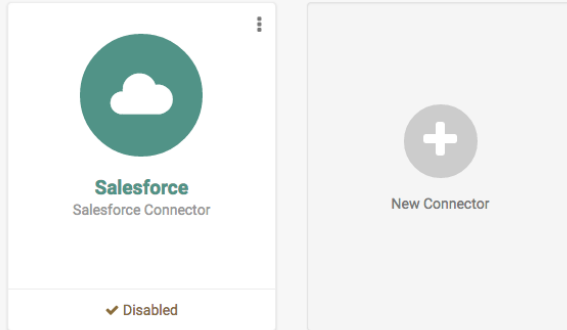
Resources

[+ New Connector](#)[+ New Managed Object](#)[+ New Mapping](#)

Connectors

[Help](#)

Connectors provide interfaces to external systems, databases, directory services, and more.



2. Enable the Salesforce connector by completing the authentication details as follows. You will need the Consumer Key and Consumer Secret that you obtained in the previous section.
 - Click on the Salesforce connector and click Edit to open the Edit Connector dialog.
 - Select True from the Enabled list.
 - Select Salesforce Connector from the Connector Type list.
 - Under Basic Connector Details select the Sandbox URL (for the purposes of testing the sample) and enter the Connector Key and Secret that you obtained in the previous section.
 - You can leave the default liveSync details for now. Click Update to update the connector configuration.

Edit Connector

General Details

Connector Name	<input type="text" value="salesforce"/>
Enabled	<input checked="" type="checkbox" value="True"/>
Connector Type	<input type="text" value="Salesforce Connector - 2.0.29.2"/>

Base Connector Details

- Your application must be configured to use this Callback URL
`https://localhost:8443/admin/oauth.html`
- You must grant this application all of the following OAuth Scopes:
 - Access your basic information
 - Access and manage your data
 - Perform requests on your behalf at any time

Production Sandbox Custom

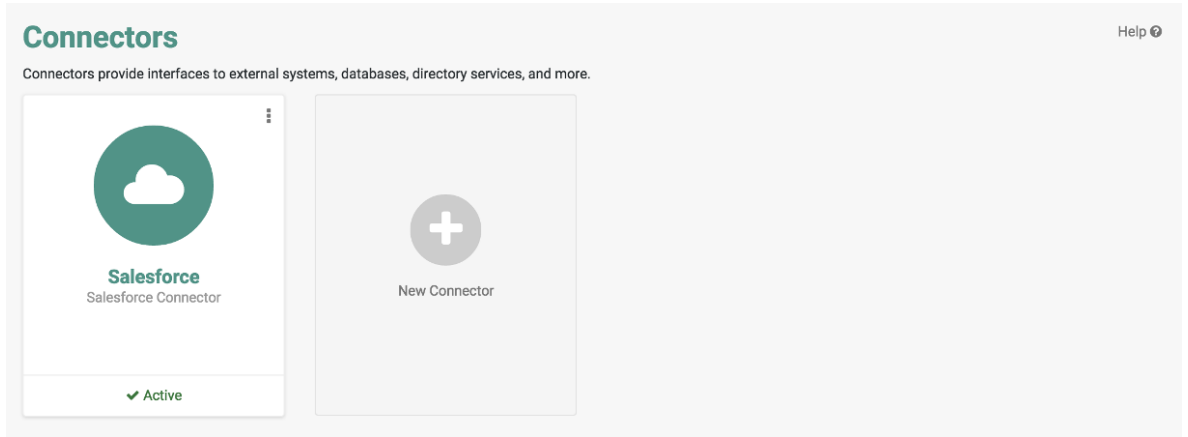
Login URL	<input type="text" value="https://login.salesforce.com/services/oauth2/token"/>
Consumer Key	<input type="text" value="3MVG9sG9Z3Q1RlbdIY.USaEytlyXe; /c_psZrWMJcT89U_JboYbruGxfoF0tbRNb.Gc4"/>
Consumer Secret	<input type="text" value="*****"/>

- On the permission request screen click Allow, to enable OpenIDM to access your Salesforce Connected App.

Note

In the current OpenIDM release, an issue is occasionally seen where the system appears to time out while retrieving the refresh token from Salesforce, at this stage. If this happens, refresh your browser and attempt the connector setup again.

On the Resources tab, your Salesforce Connector should now be Active.



- To test the reconciliation process, select the Mappings tab.

This tab shows two configured mappings, one from Salesforce to the OpenIDM repository (**managed/user**) and one from the OpenIDM repository to Salesforce.

Mapping List

+ New Mapping
Help

<div style="display: flex; align-items: center;"> <div> <p><small>SOURCE</small></p> <p>system/salesforce/User</p> <p><small>salesforce</small></p> </div> </div>	→	<div style="display: flex; align-items: center;"> <div> <p><small>TARGET</small></p> <p>managed/user</p> <p><small>managed</small></p> </div> </div>	<p>sourceSalesforceUser_managedUser</p> <p><small>Not yet reconciled.</small></p>	 Edit Delete
<div style="display: flex; align-items: center;"> <div> <p><small>SOURCE</small></p> <p>managed/user</p> <p><small>managed</small></p> </div> </div>	→	<div style="display: flex; align-items: center;"> <div> <p><small>TARGET</small></p> <p>system/salesforce/User</p> <p><small>salesforce</small></p> </div> </div>	<p>managedUser_sourceSalesforceUser</p> <p><small>Not yet reconciled.</small></p>	 Edit Delete

- Click anywhere on the first mapping and click Reconcile Now.

SOURCE **system/salesforce/User**
salesforce

TARGET **managed/user**
managed

▶ Status: Not yet reconciled.

 ⌂ Reconcile Now
 ? Help

The reconciliation operation creates the users that were present in your Salesforce organization in the OpenIDM repository.

- Retrieve the users in the repository. In the upper-right of the screen, click the `openidm-admin` link. In the pop-up menu that appears, click the Data Management View link.

This link opens the Self-Service UI. If you did not change your password in the first step, you are prompted to change your password again. You can bypass this by clicking X to close the password prompt window.

- Select the Users tab.

User List

Add User Reload Grid Clear Filters

Username	First Name	Last Name	Email	Status
admin@identityconnect-dev	Admin	User	mike.jang@forgerock.com	✓
babs.jensen@example.com	Babs	Jensen	babs.jensen@example.com	✓
cindy.venter@example.com	Cindy	Venter	cindy.venter@example.com	✓
lerato.mmutle@example.com	Lerato	Mmutle	lerato.mmutle@example.com	✓
samantha.donnely@example.com	Samantha	Donnelly	samantha.donnely@example.com	✓
steven.carter@example.com	Steven	Carter	steven.carter@example.com	✓
zawadi.treffers@example.com	Zawadi	Treffers	zawadi.treffers@example.com	✓

View 1 - 7 of 7

The users from the Salesforce organization have been reconciled to the OpenIDM repository. If the reconciliation was successful, the list of users displayed here should reflect what was in your Salesforce organization.

- To retrieve the details of a specific user, click that username on the Users tab.

The following image shows the details of user `bjensen`. Scroll down. Note the Linked Systems panel that shows the corresponding user record in Salesforce.

Babs Jensen's profile

Change Password

Username	<input type="text" value="babs.jensen@example.com"/>	✓
First Name	<input type="text" value="Babs"/>	✓
Last Name	<input type="text" value="Jensen"/>	✓
Email Address	<input type="text" value="babs.jensen@example.com"/>	✓
Role	<input type="checkbox"/> Administrator <input type="checkbox"/> Tasks Manager <input checked="" type="checkbox"/> User	
Account Status	<input type="text" value="Active"/>	⌵

- To test the second mapping (from OpenIDM to Salesforce), update any user in the OpenIDM repository. For example, update Babs Jensen's username.
- By default, *implicit synchronization* is enabled for mappings from the **managed/user** repository to any external resource. This means that when you update a managed object, any mappings defined in the `sync.json` file that have the managed object as the source are automatically executed to update the target system. For more information, see Section 14.3.2, "Mapping Source Objects to Target Objects" in the *Integrator's Guide*.

To test that the implicit synchronization has been successful, look at Babs Jensen's record in the Self-Service UI. At the bottom of the user profile, the Linked Systems panel indicates Babs Jensen's record in the Salesforce data store. Note the changed Username.

Alternatively, check the updated user record in Salesforce.

14.4. Running the Sample by Using the Command Line

Running the sample by using the command line is a little more complex. This section breaks the sample into two tasks - configuring the connector, and then testing the configuration by running reconciliation operations between the two systems.

Procedure 14.1. To Set Up the Salesforce Connector

Before you start, you will need the Consumer Key and Consumer Secret that you obtained in the previous section.

- Obtain the refresh token from salesforce.com by pointing your browser to the following URL. Substitute your Consumer Key for `CLIENT_ID`. If OpenIDM is not running on the localhost, substitute the appropriate hostname and port number in the value of the `redirect_uri` parameter.

```
https://login.salesforce.com/services/oauth2/authorize?response_type=code&
client_id=CLIENT_ID&redirect_uri=https://localhost:8443/admin/oauth.html&scope=id+api
+refresh_token
```

2. You are redirected to Salesforce, and prompted to give this application access to your Salesforce account. When you have given consent, you should receive a response URL that looks similar to the following:

```
https://localhost:8443/admin/index.html#connectors/edit//&code=aPrxJZTK7Rs03PU634VK8Jn9o_U3ZY1ERxM7Iik1F.
..
```

The `&code` part of this URL is an authorization code, that you need for the following step.

Caution

Note that this authorization code expires after 10 minutes. If you do not complete the OAuth flow within that time, you will need to start this process again.

3. Copy the authorization code from the response URL and use it as the value of the `"code"` parameter in the following REST call. You will also need to supply your Consumer Key and Consumer Secret in this call.

```
$ curl \
--verbose \
--data "grant_type=authorization_code" \
--data "client_id=consumer-key" \
--data "client_secret=consumer-secret" \
--data "redirect_uri=https://localhost:8443/admin/oauth.html" \
--data "code=access-token-code" \
"https://login.salesforce.com/services/oauth2/token"
{
  "access_token": "00DS0000003K4fU!AQMAQ0zEU
.8tCjg8Wk79yKPKCtrtaszX5jrHtoT4NBpJ8x2NFZGjg3PNuc0TWq0EgiGS_mVkfG5f4pVN5...",
  "signature": "2uREX1lseXdg3Vng/2+Hrlo/KH0WYoim+poj74wKFtw=",
  "refresh_token": "5Aep861KIwKdekr90I4iHdtDgWwRoG70_6uHrgJ
.yVtMS0UaGxRqE6WFM77w7wCV4muVMgdqKjuWI2i5S6sjN2X",
  "token_type": "Bearer",
  "instance_url": "https://example-com.cs1.my.salesforce.com",
  "scope": "id api refresh token",
  "issued_at": "1417182949781",
  "id": "https://login.salesforce.com/id/00DS0000003K4fUMAS/00530000009hWlcAAM"
}
```

The output includes an `access_token` and a `refresh_token`. You will need the `refresh_token` in the following step.

4. Edit the `configurationProperties` in your Salesforce connector configuration file (`openidm/samples/salesforce-connector/conf/provisioner.salesforce-salesforce.json`) to include your Consumer Key (`clientId`), Consumer Secret (`clientSecret`), and refresh token.


```
]
```

Procedure 14.2. Run Reconciliation by Using the Command Line

The mapping configuration file (`sync.json`) for this sample includes two mappings, `sourceSalesforceUser_managedUser`, which synchronizes users from the Salesforce with the OpenIDM repository, and `managedUser_sourceSalesforceUser`, which synchronizes changes from the OpenIDM repository to Salesforce.

1. Reconcile the repository over the REST interface by running the following command:

```
$ curl \
--cacert self-signed.crt \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request POST \
"https://localhost:8443/openidm/recon?
action=recon&mapping=sourceSalesforceUser_managedUser&waitForCompletion=true"
{
  "state": "SUCCESS",
  "_id": "8a6281ef-6faf-43dd-af5c-3a842b38c468"
}
```

The reconciliation operation returns a reconciliation run ID and the status of the operation. Reconciliation creates user objects from LDAP in the OpenIDM repository, assigning the new objects random unique IDs.

2. View the recon entry over REST for an indication of the actions that were taken on the OpenIDM repository.

```
$ curl \
--cacert self-signed.crt \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"https://localhost:8443/openidm/recon/8a6281ef-6faf-43dd-af5c-3a842b38c468"
{
  "duration": 6447,
  "ended": "2014-11-28T15:01:38.399Z",
  "started": "2014-11-28T15:01:31.952Z",
  "parameters": {
    "null": false,
    "boolean": false,
    "number": false,
    "list": false,
    "object": {
      "targetQuery": {
        "_queryId": "query-all-ids",
        "resourceName": "managed/user"
      },
      "sourceQuery": {
        "_queryId": "query-all-ids",
        "resourceName": "system/salesforce/User"
      }
    }
  }
},
```

```

"pointer": {
  "empty": true
},
"transformers": [],
"set": false,
"map": true,
"string": false,
"collection": false,
"wrappedObject": {
  "targetQuery": {
    "resourceName": "managed/user",
    "_queryId": "query-all-ids"
  },
  "sourceQuery": {
    "_queryId": "query-all-ids",
    "resourceName": "system/salesforce/User"
  }
}
},
"_id": "8a6281ef-6faf-43dd-af5c-3a842b38c468",
"mapping": "sourceSalesforceUser_managedUser",
"state": "SUCCESS",
"stage": "COMPLETED_SUCCESS",
"stageDescription": "reconciliation completed.",
"progress": {
  "links": {
    "created": 8,
    "existing": {
      "total": "0",
      "processed": 0
    }
  },
  "target": {
    "created": 8,
    "existing": {
      "total": "0",
      "processed": 0
    }
  },
  "source": {
    "existing": {
      "total": "9",
      "processed": 9
    }
  }
},
"situationSummary": {
  "FOUND_ALREADY_LINKED": 0,
  "UNASSIGNED": 0,
  "TARGET_IGNORED": 0,
  "SOURCE_IGNORED": 0,
  "MISSING": 0,
  "FOUND": 0,
  "AMBIGUOUS": 0,
  "UNQUALIFIED": 0,
  "CONFIRMED": 0,
  "SOURCE_MISSING": 0,
  "ABSENT": 9
},

```

```

"statusSummary": {
  "SUCCESS": 8,
  "FAILURE": 1
}
}

```

The output shows that eight entries were created on the target (`managed/user`).

3. You can display those users by querying the IDs in the `managed/user` repository.

```

$ curl \
  --cacert self-signed.crt \
  --header "X-OpenIDM-Username: openidm-admin" \
  --header "X-OpenIDM-Password: openidm-admin" \
  --request GET \
  "https://localhost:8443/openidm/managed/user?_queryId=query-all-ids"
{
  "remainingPagedResults": -1,
  "pagedResultsCookie": null,
  "resultCount": 8,
  "result": [
    {
      "_rev": "0",
      "_id": "f15322f2-5873-4e5f-a4e5-2d4bc03dd190"
    },
    {
      "_rev": "0",
      "_id": "85879c60-afa1-4425-8c7a-5cccbbaff587"
    },
    {
      "_rev": "0",
      "_id": "ed3fe655-29a6-4016-b6bc-4b2356911fd1"
    },
    {
      "_rev": "0",
      "_id": "34678464-c080-41b1-8da6-d5fde9d35aeb"
    },
    {
      "_rev": "0",
      "_id": "02d5da29-8349-4f35-affc-5f6c331307ef"
    },
    {
      "_rev": "0",
      "_id": "f91d6fce-bf27-4379-9411-fd626f8a9528"
    },
    {
      "_rev": "0",
      "_id": "6ace9220-59e7-4d97-8683-e03362a9150c"
    },
    {
      "_rev": "0",
      "_id": "56863eea-35d7-4aeb-a017-74ef28fd3116"
    }
  ]
}

```


Chapter 15

Scripted Kerberos Connector Sample

New in OpenIDM 4.5.0, the scripted Kerberos connector sample demonstrates how to manage Kerberos user principals and how to reconcile user principals with OpenIDM managed user objects.

The connector configuration (`/path/to/openidm/samples/kerberos/conf/provisioner.openicf-kerberos.json`) assumes that OpenIDM is running on a host that is separate from the Kerberos host.

This sample assumes that the default realm is `EXAMPLE.COM` and that there is an existing user principal `openidm/admin`. Adjust the sample to match your Kerberos realm and principals.

15.1. Editing the Kerberos Connector Configuration

Before you run this sample, edit the connector configuration file to match your Kerberos environment. Specifically, set the correct values for the following properties:

host

The host name or IP address of the machine on which Kerberos is running.

port

The SSH port on that machine.

Default: `22` (the default SSH port)

user

The username of the account that is used to connect to the SSH server.

password

The password of the account that is used to connect to the SSH server.

prompt

A string that represents the remote SSH session prompt. This must be the exact prompt string, in the format `username@target:`, for example `root@localhost:~$`. The easiest way to obtain this string is to `ssh` into the machine and copy paste the prompt.

customConfiguration

The details of the admin user principal and the default realm.

This example assumes an admin user principal of `openidm/admin`.

For more information on setting this property, see [customConfiguration](#) in the *Connectors Guide*.

customSensitiveConfiguration

The password for the user principal.

For more information on setting this property, see [customSensitiveConfiguration](#) in the *Connectors Guide*.

Your connector configuration should look something like the following:

```
...
"configurationProperties" : {
  "host" : "192.0.2.0",
  "port" : 22,
  "user" : "admin",
  "password" : "Passw0rd",
  "prompt" : "admin@myhost:~$",
  "sudoCommand" : "/usr/bin/sudo",
  "echoOff" : true,
  "terminalType" : "vt102",
  "setLocale" : false,
  "locale" : "en_US.utf8",
  "connectionTimeout" : 5000,
  "expectTimeout" : 5000,
  "authenticationType" : "PASSWORD",
  "throwOperationTimeoutException" : true,
  "customConfiguration" : "kadmin { cmd = '/usr/sbin/kadmin.local'; user='openid/admin';
default_realm='EXAMPLE.COM' }",
  "customSensitiveConfiguration" : "kadmin { password = 'Passw0rd'}",
...

```

OpenIDM will encrypt all passwords in the configuration when it starts up, or whenever it reloads the configuration file.

For information about the complete Kerberos connector configuration, see Section 10.2, "Configuring the Kerberos Connector" in the *Connectors Guide*.

Caution

Do not modify the value of the `scriptRoots` or `classpath` properties unless you have extracted the scripts from the connector bundle and placed them on the filesystem.

15.2. Running the Kerberos Sample

The commands in this section achieve the following:

1. Start OpenIDM and check that the connector can reach the Kerberos server.
2. Create two users in the OpenIDM managed repository.
3. Reconcile the managed repository with the Kerberos server so that the new users are created in Kerberos.

4. Retrieve the details of one of the new Kerberos principals from the server.
5. Delete one of the managed users.
6. Reconcile the managed repository again and note that the corresponding Kerberos principal has been deleted.
1. Start OpenIDM with the configuration for the Kerberos sample:

```
$ cd /path/to/openidm
```

```
$ startup.sh -p samples/kerberos
```

2. Test that your connector configuration is correct and that OpenIDM can reach your Kerberos server, with the following command:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request POST \
"http://localhost:8080/openidm/system?_action=test"
[
  {
    "name": "kerberos",
    "enabled": true,
    "config": "config/provisioner.openicf/kerberos",
    "objectTypes": [
      "_ALL_",
      "account"
    ],
  },
  "connectorRef": {
    "bundleName": "org.forgerock.openicf.connectors.kerberos-connector",
    "connectorName": "org.forgerock.openicf.connectors.kerberos.KerberosConnector",
    "bundleVersion": "[1.4.0.0,1.5.0.0)"
  },
  "displayName": "Kerberos Connector",
  "ok": true
}
]
```

If the command returns `"ok": true`, as in the preceding output, your configuration is correct and you can continue with the sample.

3. Retrieve a list of the existing user principals in the Kerberos database:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/system/kerberos/account?_queryId=query-all-ids"
{
  "result": [
    {
      "_id": "K/M@EXAMPLE.COM",
      "principal": "K/M@EXAMPLE.COM"
    },
    {
```

```

    "_id": "kadmin/admin@EXAMPLE.COM",
    "principal": "kadmin/admin@EXAMPLE.COM"
  },
  {
    "_id": "kadmin/changepw@EXAMPLE.COM",
    "principal": "kadmin/changepw@EXAMPLE.COM"
  },
  {
    "_id": "kadmin/krbl.example.com@EXAMPLE.COM",
    "principal": "kadmin/krbl.example.com@EXAMPLE.COM"
  },
  {
    "_id": "kiprop/krbl.example.com@EXAMPLE.COM",
    "principal": "kiprop/krbl.example.com@EXAMPLE.COM"
  },
  {
    "_id": "krbtgt/EXAMPLE.COM@EXAMPLE.COM",
    "principal": "krbtgt/EXAMPLE.COM@EXAMPLE.COM"
  },
  {
    "_id": "openidm/admin@EXAMPLE.COM",
    "principal": "openidm/admin@EXAMPLE.COM"
  }
],
...
}

```

4. Create two new managed users, either over REST or by using the Admin UI.

The following command creates users bjensen and scarter over REST. To create similar users by using the Admin UI, select Managed > User and click New User:

```

$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--header "Content-type: application/json" \
--request POST \
--data '{
  "userName": "bjensen",
  "givenName": "Barbara",
  "sn": "Jensen",
  "password": "Passw0rd",
  "displayName": "Barbara Jensen",
  "mail": "bjensen@example.com"
}' \
"http://localhost:8080/openidm/managed/user?_action=create"
{
  "_id": "ce3d9b8f-1d15-4950-82c1-f87596aadcb6",
  "_rev": "2",
  "userName": "bjensen",
  "givenName": "Barbara",
  "sn": "Jensen",
  "displayName": "Barbara Jensen",
  "mail": "bjensen@example.com",
  "accountStatus": "active",
  "effectiveRoles": [],
  "effectiveAssignments": []
}

```

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--header "Content-type: application/json" \
--request POST \
--data '{
  "userName": "scarter",
  "givenName": "Steven",
  "sn" : "Carter",
  "password" : "Passw0rd",
  "displayName": "Steven Carter",
  "mail" : "scarter@example.com"
}' \
"http://localhost:8080/openidm/managed/user?_action=create"
{
  "_id": "a204ca60-b0fc-42f8-bf93-65bb30131361",
  "_rev": "2",
  "userName": "scarter",
  "givenName": "Steven",
  "sn": "Carter",
  "displayName": "Steven Carter",
  "mail": "scarter@example.com",
  "accountStatus": "active",
  "effectiveRoles": [],
  "effectiveAssignments": []
}
```

5. Run a reconciliation operation between the managed user repository and the Kerberos database to create the new users bjensen and scarter in Kerberos. You can run the reconciliation over REST, or using the Admin UI.

The following command creates runs the reconciliation over REST:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request POST \
"http://localhost:8080/openidm/recon?_action=recon&mapping=managedUser_systemKerberos"
{
  "_id": "862ab9ba-d1d9-4058-b6bc-a23a94b68776-234",
  "state": "ACTIVE"
}
```

To run the reconciliation by using the Admin UI, select Configure > Mappings, click on the [managedUser_systemKerberos](#) mapping, and click Reconcile Now.

6. Retrieve the list of Kerberos user principals again. You should now see bjensen and scarter in this list:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/system/kerberos/account?_queryId=query-all-ids"
{
  "result": [
    {
      "_id": "bjensen@EXAMPLE.COM",
      "principal": "bjensen@EXAMPLE.COM"
    },
    {
      "_id": "scarter@EXAMPLE.COM",
      "principal": "scarter@EXAMPLE.COM"
    },
    ...
    {
      "_id": "openidm/admin@EXAMPLE.COM",
      "principal": "openidm/admin@EXAMPLE.COM"
    }
  ],
  ...
}
```

7. Retrieve bjensen's complete user principal from the Kerberos server:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/system/kerberos/account/bjensen@EXAMPLE.COM"
{
  "_id": "bjensen@EXAMPLE.COM",
  "lastFailedAuthentication": "[never]",
  "passwordExpiration": "[none]",
  "lastSuccessfulAuthentication": "[never]",
  "maximumTicketLife": "0 days 10:00:00",
  "lastModified": "Tue May 24 04:05:45 EDT 2016 (openidm/admin@EXAMPLE.COM)",
  "policy": "user [does not exist]",
  "expirationDate": "[never]",
  "failedPasswordAttempts": "0",
  "maximumRenewableLife": "7 days 00:00:00",
  "principal": "bjensen@EXAMPLE.COM",
  "lastPasswordChange": "Tue May 24 04:05:45 EDT 2016"
}
```

Note the default values for properties such as `maximumRenewableLife`. These values are set in your connector configuration. For more information, see Section 10.2, "Configuring the Kerberos Connector" in the *Connectors Guide*.

To perform this step in the Admin UI, select Manage > User, click bjensen's entry, and click the Linked Systems tab to display her corresponding entry on the Kerberos server.

8. Delete the managed user bjensen by specifying her managed object ID in the DELETE request.

First, obtain her ID by querying for her `userName`:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/managed/user?_queryFilter=username+eq+'bjensen'"
{
  "result": [
    {
      "_id": "ce3d9b8f-1d15-4950-82c1-f87596aadcb6",
      "_rev": "3",
      "userName": "bjensen",
      "givenName": "Barbara",
      "sn": "Jensen",
      "displayName": "Barbara Jensen",
      "mail": "bjensen@example.com",
      "accountStatus": "active",
      "effectiveRoles": [],
      "effectiveAssignments": []
    }
  ],
  ...
}
```

Now delete the user with ID `ce3d9b8f-1d15-4950-82c1-f87596aadcb6`. This ID will obviously be different in your example.

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request DELETE \
"http://localhost:8080/openidm/managed/user/ce3d9b8f-1d15-4950-82c1-f87596aadcb6"
{
  "_id": "ce3d9b8f-1d15-4950-82c1-f87596aadcb6",
  "_rev": "3",
  "userName": "bjensen",
  "givenName": "Barbara",
  "sn": "Jensen",
  "displayName": "Barbara Jensen",
  "mail": "bjensen@example.com",
  "accountStatus": "active",
  "effectiveRoles": [],
  "effectiveAssignments": []
}
```

To delete bjensen's managed user entry by using the Admin UI, select Manage > User, click on bjensen's entry, select the checkbox next to her entry, and click Delete Selected.

9. Reconcile the managed user repository and the Kerberos database again:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request POST \
"http://localhost:8080/openidm/recon?_action=recon&mapping=managedUser_systemKerberos"
{
  "_id": "862ab9ba-d1d9-4058-b6bc-a23a94b68776-584",
  "state": "ACTIVE"
}
```

10. Retrieve the list of Kerberos user principals again. The Kerberos principal for bjensen should have been removed from the list:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/system/kerberos/account?_queryId=query-all-ids"
{
  "result": [
    {
      "_id": "K/M@EXAMPLE.COM",
      "principal": "K/M@EXAMPLE.COM"
    },
    {
      "_id": "kadmin/admin@EXAMPLE.COM",
      "principal": "kadmin/admin@EXAMPLE.COM"
    },
    {
      "_id": "kadmin/changepw@EXAMPLE.COM",
      "principal": "kadmin/changepw@EXAMPLE.COM"
    },
    {
      "_id": "kadmin/krbl.example.com@EXAMPLE.COM",
      "principal": "kadmin/krbl.example.com@EXAMPLE.COM"
    },
    {
      "_id": "kiprop/krbl.example.com@EXAMPLE.COM",
      "principal": "kiprop/krbl.example.com@EXAMPLE.COM"
    },
    {
      "_id": "krbtgt/EXAMPLE.COM@EXAMPLE.COM",
      "principal": "krbtgt/EXAMPLE.COM@EXAMPLE.COM"
    },
    {
      "_id": "scarter@EXAMPLE.COM",
      "principal": "scarter@EXAMPLE.COM"
    },
    {
      "_id": "openidm/admin@EXAMPLE.COM",
      "principal": "openidm/admin@EXAMPLE.COM"
    }
  ],
  ...
}
```


Note

Some user IDs in Kerberos include characters such as a forward slash (/) and an "at sign" (@) that prevent them from being used directly in a REST URL. For example, `openidm/system/kerberos/account/kadmin/admin@EXAMPLE.COM`, where the ID is `kadmin/admin@EXAMPLE.COM`. To retrieve such entries directly over REST, you must URL-encode the Kerberos ID as follows:

```
"http://localhost:8080/openidm/system/kerberos/account/kadmin%2Fadmin%40EXAMPLE.COM"
```

Chapter 16

Customer Data Management Sample

This sample demonstrates the CDM process to register and authenticate users with multiple social identity providers - Google, Facebook, and LinkedIn. The sample also demonstrates the use of the Marketo connector to reconcile this data to a Marketo database for lead generation (marketing).

This sample assumes that you have a Marketo administrative account, with privileges to create a custom service. From that service, you'll need Client ID and Client Secret information, among other things. For more information, see Section 16.2, "Setting Up Users for Marketo".

Think of this sample as a practical, "hands-on" HOWTO; for detailed configuration options, see the following documentation:

- Chapter 10, "Configuring Social ID Providers" in the *Integrator's Guide*.
- Chapter 12, "Marketo Connector" in the *Connectors Guide*.

16.1. Set Up Registration and Authentication

This sample includes files with a pre-configured set of Social ID Providers. It also includes a `SOCIAL_PROVIDERS` authentication module that works with the specified social identity providers. In this sample, you'll enable those social identity providers for user authentication and registration.

In the following procedure, you'll add your own Client ID and Client Secret values before activating each social ID provider. If you do not already have these values, refer to the following instructions:

- Section 10.3.1, "Setting Up Google" in the *Integrator's Guide*.
- Section 10.4.1, "Setting Up LinkedIn" in the *Integrator's Guide*.
- Section 10.5.1, "Setting Up Facebook" in the *Integrator's Guide*. When you configure a Facebook app, look for the values for App ID and App Secret.

The purpose of this procedure is to demonstrate how you can start using social ID providers as quickly as possible, through the Admin UI.

1. Log in to the Admin UI at `https://localhost:8443/admin` as the default administrative user `openidm-admin` with password `openidm-admin`.
2. Select Configure > Social ID Providers. Highlight each provider. Add values for Client ID and Client Secret. Toggle to activate the Enabled switch, and then select Save.

3. Social ID provider-based user registration should already be enabled by default. To confirm, select Configure > User Registration.

You can review the configuration in the following JSON files in the `openidm/samples/cdm/conf` directory:

- `identityProvider-google.json`
- `identityProvider-facebook.json`
- `identityProvider-linkedIn.json`
- `authentication.json`
- `selfservice-registration.json`
- `ui-configuration.json`

Once you've configured one or more social ID providers, you can use a social ID account to register with OpenIDM. After registration, your account will be displayed in the OpenIDM managed user repository.

For details, see Chapter 10, "*Configuring Social ID Providers*" in the *Integrator's Guide*.

16.2. Setting Up Users for Marketo

This sample reconciles all qualifying OpenIDM users, including those registered through a social ID provider, to the Marketo database.

Review the mapping. Navigate to Configure > Mappings, and select the `managed/user` to `system/marketo/account` mapping.

Under the Association tab, select Individual Record Validation. The Valid Source condition, `Send me special offers and services`, limits reconciliation to users who have accepted that condition.

MAPPING DETAIL managedUser_systemMarketoAccount Reconcile

SOURCE managed/user managed

TARGET system/marketo/account marketo

► Status: Not yet reconciled. Help

Properties **Association** Behaviors Scheduling Advanced

► **Reconciliation Query Filters**

▼ **Individual Record Validation**

Define valid source and target records for reconciliation

Valid Source Validate based on user preferences

Reconcile records with these user preferences

- Send me news and updates
- Send me special offers and services

Valid Target All records valid

Save

When a user registers with ForgeRock IDM, they can choose to accept the noted condition. As a regular user, they can also select (or deselect) the noted condition in the self-service UI. To do so, log into ForgeRock IDM at <http://localhost:8080/>, and select Profile. On the User Profile page select the preferences tab.

For more information on how preferences work in a mapping, see Section 14.3.4, "Configuring Synchronization Filters With User Preferences" in the *Integrator's Guide*.

16.3. Configuring the Marketo Connector

Before going through the following procedure, make sure you have values for Client ID and Client Secret for your Marketo service. For the Marketo procedure, see their *Quick Start Guide for Marketo REST API*.

1. Open the Marketo connector. In the Admin UI, select Configure > Connectors, and select the connector.
2. Include appropriate information for the fields shown in the following list:
 - Instance: Find this FQDN in within the REST API endpoint associated with your Marketo web service. It may be a value such as *some-number.mktorest.com*.
 - Client ID: Find this in the details of your Marketo service *LaunchPoint*.
 - Client Secret: Find this in the details of your Marketo service *LaunchPoint*.
 - List Name: If you have created a Marketo Group Name, enter it here.
 - Lead Fields: Normally left blank; refer to Marketo documentation for options.
 - Partition Name: may be left blank. Standard Marketo accounts do not support partitions.

For more information on these fields, see *Chapter 12, "Marketo Connector"* in the *Connectors Guide*.

As a Marketo administrator, you should be able to find values for each of these fields for your Marketo custom service.
3. When you've finished making changes, press Save.

16.4. Reviewing Marketo Leads

After ForgeRock Identity Management reconciles data to the Marketo database, you should have a list of users who have accepted the rules associated with Individual Record Validation. You should find those users in the Marketo group list that you configured with your Marketo administrative account.

If any of your users deselect their applicable marketing preferences, as described in Section 16.2, "Setting Up Users for Marketo", ForgeRock IDM removes those accounts from the Marketo database upon the next reconciliation.

Chapter 17

Custom Endpoint Sample

This chapter describes the custom endpoint sample delivered with OpenIDM.

OpenIDM supports scriptable custom endpoints that enable you to launch arbitrary scripts through an OpenIDM REST URI. For information about how custom endpoints are configured, see Section 15.2, "Creating Custom Endpoints to Launch Scripts" in the *Integrator's Guide*.

The sample endpoint provided in `/path/to/openidm/samples/customendpoint` illustrates the configuration of a custom endpoint, and the structure of custom endpoint scripts.

The purpose of this custom endpoint is to return a list of variables available to each method used in a script. The scripts show the complete set of methods that can be used. These methods map to the standard HTTP verbs - create, read, update, delete, patch, query, and action. A sample JavaScript and Groovy script is provided.

To run the sample:

1. Copy the endpoint configuration file (`samples/customendpoint/conf/endpoint-echo.json`) to your project's `conf` directory.
2. Copy either the JavaScript file (`samples/customendpoint/script/echo.js`) or Groovy script file (`samples/customendpoint/script/echo.groovy`) to your project's `script` directory.
3. Open the endpoint configuration file in a text editor:

```
{
  "file" : "echo.groovy",
  "type" : "groovy",
  "_file" : "echo.js",
  "_type" : "text/javascript"
}
```

The configuration file contains nothing more than a reference to the endpoint scripts. In this case, the JavaScript script is commented out (with an underscore before the `file` and `type` properties. If you want to use the JavaScript endpoint script, uncomment these lines and comment out the lines that correspond to the Groovy script in the same way.

Endpoint configuration files can include a `context` property that specifies the route to the endpoint, for example:

```
"context" : "endpoint/linkedView/*"
```

If no `context` is specified, the route to the endpoint is taken from the file name, in this case `endpoint/echo`.

4. Test each method in succession to return the expected request structure of that method. The following examples show the request structure of the read, create and patch methods. The configuration file has been edited to use the JavaScript file, rather than the Groovy file. The output shown in these examples has been cropped for legibility. For a description of each parameter, see Section 15.2.2, "Writing Custom Endpoint Scripts" in the *Integrator's Guide*.

The following command performs a read on the echo endpoint and returns the request structure of a read request:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/endpoint/echo"
{
  "_id": "",
  "method": "read",
  "resourceName": "",
  "parameters": {},
  "context": {
    ...
  }
}
```

The following command performs a query on the echo endpoint and returns the request structure of a query request:

```
$ curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--request GET \
"http://localhost:8080/openidm/endpoint/echo?_queryId=query-all-ids"
{
  "result": [
    {
      "method": "query",
      "resourceName": "",
      "pagedResultsCookie": null,
      "pagedResultsOffset": 0,
      "pageSize": 0,
      "queryExpression": null,
      "queryId": "query-all-ids",
      "queryFilter": "null",
      "parameters": {},
      "content": null,
      "context": {
        ...
      }
    }
  ],
  ...
}
```

The following command sends a create request to the echo endpoint. No user is actually created. The endpoint script merely returns the request structure of a create request. The **content** parameter in this case provides the JSON object that was sent with the request:

```
$ curl \
--header "X-OpenIDM-Password: openidm-admin" \
--header "X-OpenIDM-Username: openidm-admin" \
--header "Content-Type: application/json" \
--data '{
  "userName": "steve",
  "givenName": "Steve",
  "sn": "Carter",
  "telephoneNumber": "0828290289",
  "mail": "scarter@example.com",
  "password": "Passw0rd"
}' \
--request POST \
"http://localhost:8080/openidm/endpoint/echo?action=create"
{
  "_id": "",
  "method": "create",
  "resourceName": "",
  "newResourceId": null,
  "parameters": {},
  "content": {
    "userName": "steve",
    "givenName": "Steve",
    "sn": "Carter",
    "telephoneNumber": "0828290289",
    "mail": "scarter@example.com",
    "password": "Passw0rd"
  },
  "context": {
    ...
  }
}
```

The following command sends a patch request to the echo endpoint.


```
$ curl \
--header "X-OpenIDM-Password: openidm-admin" \
--header "X-OpenIDM-Username: openidm-admin" \
--header "Content-Type: application/json" \
--data '[
  {
    "operation": "replace",
    "field": "/givenName",
    "value": "Steven"
  }
]' \
--request PATCH \
"http://localhost:8080/openidm/endpoint/echo"
{
  "_id": "",
  "method": "patch",
  "resourceName": "",
  "revision": null,
  "patch": [
    {
      "operation": "replace",
      "field": "/givenName",
      "value": "Steven"
    }
  ],
  "parameters": {},
  "context": {
    ...
  }
}
```

Glossary

correlation query	<p>A correlation query specifies an expression that matches existing entries in a source repository to one or more entries on a target repository. While a correlation query may be built with a script, it is <i>not</i> a correlation script.</p> <p>As noted in Section 14.3.2.6, "Correlating Source Objects With Existing Target Objects" in the <i>Integrator's Guide</i>, you can set up a query definition, such as <code>_queryId</code>, <code>_queryFilter</code>, or <code>_queryExpression</code>, possibly with the help of a <code>linkQualifier</code>.</p>
correlation script	<p>A correlation script matches existing entries in a source repository, and returns the IDs of one or more matching entries on a target repository. While it skips the intermediate step associated with a <code>correlation query</code>, a correlation script can be relatively complex, based on the operations of the script.</p>
entitlement	<p>An entitlement is a collection of attributes that can be added to a user entry via roles. As such, it is a specialized type of <code>assignment</code>. A user or device with an entitlement gets access rights to specified resources. An entitlement is a property of a managed object.</p>
JSON	<p>JavaScript Object Notation, a lightweight data interchange format based on a subset of JavaScript syntax. For more information, see the JSON site.</p>
JWT	<p>JSON Web Token. As noted in the <i>JSON Web Token draft IETF Memo</i>, "JSON Web Token (JWT) is a compact URL-safe means of representing claims to be transferred between two parties." For OpenIDM, the JWT is associated with the <code>JWT_SESSION</code> authentication module.</p>

managed object	An object that represents the identity-related data managed by OpenIDM. Managed objects are configurable, JSON-based data structures that OpenIDM stores in its pluggable repository. The default configuration of a managed object is that of a user, but you can define any kind of managed object, for example, groups or roles.
mapping	A policy that is defined between a source object and a target object during reconciliation or synchronization. A mapping can also define a trigger for validation, customization, filtering, and transformation of source and target objects.
OSGi	A module system and service platform for the Java programming language that implements a complete and dynamic component model. For a good introduction, see the OSGi site . While OpenIDM services are designed to run in any OSGi container, currently only Apache Felix is supported.
reconciliation	During reconciliation, comparisons are made between managed objects and objects on source or target systems. Reconciliation can result in one or more specified actions, including, but not limited to, synchronization.
resource	An external system, database, directory server, or other source of identity data to be managed and audited by the identity management system.
REST	Representational State Transfer. A software architecture style for exposing resources, using the technologies and protocols of the World Wide Web. REST describes how distributed data objects, or resources, can be defined and addressed.
role	OpenIDM includes two different types of provisioning roles and authorization roles. For more information, see Section 9.4, "Working With Managed Roles" in the <i>Integrator's Guide</i> .
source object	In the context of reconciliation, a source object is a data object on the source system, that OpenIDM scans before attempting to find a corresponding object on the target system. Depending on the defined mapping, OpenIDM then adjusts the object on the target system (target object).
synchronization	The synchronization process creates, updates, or deletes objects on a target system, based on the defined mappings from the source system. Synchronization can be scheduled or on demand.
system object	A pluggable representation of an object on an external system. For example, a user entry that is stored in an external LDAP directory is represented as a system object in OpenIDM for the period during which OpenIDM requires access to that entry. System objects follow

the same RESTful resource-based design principles as managed objects.

target object

In the context of reconciliation, a target object is a data object on the target system, that OpenIDM scans after locating its corresponding object on the source system. Depending on the defined mapping, OpenIDM then adjusts the target object to match the corresponding source object.

Index

A

Authentication Module
ForgeRock Identity Provider, 220

I

Installing
Samples, 6

M

Messaging, 149

O

OpenAM
OpenID Connect
Options, 217

S

Samples
Google Apps Connector, 250
Salesforce Connector, 261
Sample 2 - LDAP one way, 23
Sample 2b - LDAP two way, 27
Sample 2c - Synchronizing LDAP Group
Membership, 32
Sample 2d - Synchronizing LDAP Groups, 38
Sample 5 - Synchronization of two resources,
42
Sample 5b - Failure Compensation With
Multiple Resources, 45
Sample 6 - LiveSync between two LDAP
servers, 48
ScriptedCREST Connector, 99
ScriptedREST Connector, 89