



Upgrade Guide

/ ForgeRock Identity Management 7.1

Latest update: 7.1.0

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

Copyright © 2011-2021 ForgeRock AS.

Abstract

This guide shows you how to upgrade an existing deployment to the latest ForgeRock® Identity Management release.



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.

© Copyright 2010-2020 ForgeRock, Inc. All rights reserved. ForgeRock is a registered trademark of ForgeRock, Inc. Other marks appearing herein may be trademarks of their respective owners.

This product or document is protected by copyright and distributed under licenses restricting its use, copying, and distribution. No part of this product or document may be reproduced in any form by any means without prior written authorization of ForgeRock and its licensors, if any.

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESSED OR IMPLIED CONDITIONS, REPRESENTATIONS, AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

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 NON-INFRINGEMENT 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@gnome.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 NON-INFRINGEMENT 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@free.fr.

FontAwesome Copyright

Copyright (c) 2017 by Dave Gandy, <https://fontawesome.com/>.

This Font Software is licensed under the SIL Open Font License, Version 1.1. See <https://opensource.org/licenses/OFL-1.1>.




Table of Contents

Overview	iv
1. About Upgrades	1
Supported Upgrade Paths	1
2. Before You Upgrade	2
3. Place a Server in Maintenance Mode	4
4. Migrate Your Configuration	6
Migrate Configuration Files	6
Migrate <code>boot.properties</code>	7
Migrate Security Settings	7
Migrate Custom Scripts	7
Migrate Custom Bundles	8
Migrate Provisioner Files	8
Migrate UI Customizations	8
5. Update the Repository	9
Upgrade an Existing Repository	9
Create a New Repository	11
6. Migrate Data	12
Configure the Migration Service	13
Run the Data Migration	17
7. Upgrade a Clustered Deployment	19
IDM Glossary	20

Overview

This guide shows you how to upgrade an existing deployment to the latest ForgeRock Identity Management release.

Quick Start

 <p>Migrate Configuration</p> <p>Migrate an existing IDM configuration to IDM 7.1.</p>	 <p>Update Repository</p> <p>Update an existing repository or install a new repository for IDM 7.1.</p>	 <p>Migrate Data</p> <p>Move the data in an existing IDM repository to an updated deployment.</p>
---	--	--

The upgrade process is largely dependent on your deployment and on the extent to which you have customized IDM. Engage [ForgeRock Support Services](#) for help in upgrading an existing deployment. Also, read the [Release Notes](#) before you start an upgrade; specifically, "*Incompatible Changes*".

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

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

Chapter 1

About Upgrades

The automated update process available with previous IDM versions is no longer supported. This chapter describes the manual process required to upgrade an existing IDM deployment. At a high level, the manual update process involves the following steps:

1. Install IDM 7.1.
2. Optionally, place your existing server in maintenance mode.
3. Migrate your existing IDM configuration to the new installation.
4. Update your repository.
5. Test that your scripts and customizations work as expected.
6. Migrate existing data to the new installation.

Supported Upgrade Paths

The following table contains information about the supported upgrade paths to IDM 7.1:

Upgrade Paths

Version	Upgrade Supported to IDM 7.1
IDM 7.0.x	✓
IDM 6.5.x	✓
IDM 6.0.x	✓
IDM 5.5.x	✓
IDM 5.0.x	✓

Important

Depending on how you have customized your deployment, there might be incompatible configuration changes when you upgrade from versions prior to IDM 6.5.x. Read the upgrade documentation for each interim release and apply all required script and configuration changes.

Chapter 2

Before You Upgrade

Fulfill these requirements before you upgrade IDM, especially before upgrading the software in a production environment. Also see the requirements listed in "*Before You Install*" and the changes listed in "*Incompatible Changes*" in the *Release Notes*.

Before you start, verify that you have a supported Java version installed:

Supported Java Versions

Vendor	Versions
OpenJDK, including OpenJDK-based distributions: <ul style="list-style-type: none">• AdoptOpenJDK/Eclipse Adoptium• Amazon Corretto• Azul Zulu• Red Hat OpenJDK ForgeRock tests most extensively with AdoptOpenJDK/Eclipse Adoptium. ForgeRock recommends using the HotSpot JVM.	11
Oracle Java	11

If the server uses an older version that is no longer supported, install a newer Java version before you update, and follow the instructions in "*Check Your Java Installation*" in the *Installation Guide*.

Then, follow these steps:

1. Back up your existing deployment by archiving the `openidm` directory and creating a backup of the repository and all other applicable databases.

Note

If you use workflow, you must manually dump the workflow database tables, and then import them *before* you start the new instance of IDM for the first time. The workflow database tables start with the prefix `ACT_`. For information on how to dump/import individual tables, see the documentation for your database.

2. To save a record of the audit logs from your existing IDM installation, manually copy the log files from the `/path/to/openidm/audit/` directory, before you start the upgrade.
3. Download and extract `IDM-7.1.0.zip` from the ForgeRock BackStage download site.

Chapter 3

Place a Server in Maintenance Mode

The maintenance service disables non-essential services of a running IDM instance, in preparation for an update to a later version. When maintenance mode is enabled, services such as recon, sync, scheduling, and workflow are disabled. The complete list of disabled services is output to the log file.

The router remains functional and requests to the `maintenance` endpoint continue to be serviced. Requests to endpoints that are serviced by a disabled component return the following response:

```
404 Resource endpoint-name not found
```

Before you enable maintenance mode, temporarily suspend scheduled tasks.

Manage maintenance mode over REST as follows:

+ *Enable maintenance mode*

```
curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--header "Accept-API-Version: resource=1.0" \
--request POST \
"http://localhost:8080/openidm/maintenance?_action=enable"
{
  "maintenanceEnabled": true
}
```

+ *Disable maintenance mode*

```
curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--header "Accept-API-Version: resource=1.0" \
--request POST \
"http://localhost:8080/openidm/maintenance?_action=disable"
{
  "maintenanceEnabled": false
}
```

+ *Check if a server is in maintenance mode*


```
curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--header "Accept-API-Version: resource=1.0" \
--request POST \
"http://localhost:8080/openidm/maintenance?_action=status"
{
  "maintenanceEnabled": false
}
```

If the server is in maintenance mode, the command returns `"maintenanceEnabled": true`; otherwise it returns `"maintenanceEnabled": false`.

Chapter 4

Migrate Your Configuration

This chapter covers the steps required to migrate your IDM configuration to IDM 7.1.

There is no automated way to migrate a customized configuration to IDM 7.1, so you must migrate customized configuration files manually. Assuming you are upgrading from IDM 7.0.x, there are three ways to do this:

- Use the new IDM 7.1 configuration files as a base, and copy any customizations you have made to the new files.

This is the preferred option, particularly if you have used version control on your configuration and can determine the exact changes you have applied.

- Use your existing configuration files as a base, and add any new IDM 7.1 configuration to your existing files.
- Use your existing configuration "as is" with no IDM 7.1 changes.

In most cases, a customized IDM 7.0.x configuration will work without further modification on IDM 7.1.

- "Migrate Configuration Files"
- "Migrate `boot.properties`"
- "Migrate Security Settings"
- "Migrate Custom Scripts"
- "Migrate Custom Bundles"
- "Migrate Provisioner Files"
- "Migrate UI Customizations"

Migrate Configuration Files

For customized files in your project's `conf/` directory, check that the customizations are compatible with the changes outlined in "*Incompatible Changes*" in the *Release Notes*. If there are no incompatible changes, either copy your old configuration files to your IDM 7.1 installation, or copy any customization into the corresponding new configuration files.

Migrate `boot.properties`

On the IDM 7.1 installation, edit the `resolver/boot.properties` file to match any customizations that you made on your IDM 7.0.x server. Specifically, check the following elements:

- The HTTP, HTTPS, and mutual authentication ports.

If you changed the default ports in your IDM 7.0.x deployment, make those same changes in the new `boot.properties` file.

- Check that the keystore and truststore passwords match the current passwords for the keystore and truststore of your existing IDM deployment.

Migrate Security Settings

Copy the contents of your IDM 7.0.x `security/` folder to the IDM 7.1 installation.

Warning

If you do not copy your old truststore and keystore files to your new instance, you will be unable to decrypt anything that was encrypted by your old instance of IDM.

Migrate Custom Scripts

Migrate any custom scripts or default scripts *that you have modified* to the `script` directory of your IDM 7.1 instance. In general, custom and customized scripts should be located in the `openidm/script` directory of your existing IDM deployment.

For custom scripts, review "*Incompatible Changes*" in the *Release Notes*. If you are confident that the scripts will work as intended on IDM 7.1, copy these scripts to the new instance.

If you modified a default IDM script, compare the default versions of the IDM 7.0.x and IDM 7.1 scripts. If nothing has changed between the default versions, review your customizations against "*Incompatible Changes*" in the *Release Notes*. If you are confident that your changes will work as intended on the new version, copy the customized scripts to the new `script` directory.

If a default script has changed since the IDM 7.0.x release, test that your customizations work with the new default script before porting your changes to that new script. The following default scripts have changed in IDM 7.1:

- `policy.js`
- `router-authz.js`
- `propertiesModifiedFilter.groovy`

Note

If you modify any shell scripts, such as `startup.sh`, you must migrate your changes manually to the new version of the script.

Migrate Custom Bundles

If your existing deployment includes any custom JAR files in the `bundles` directory, migrate these to the new deployment. Pay particular attention to any files that support JDBC database drivers.

Migrate Provisioner Files

Change any customized provisioner configurations in your existing deployment to point to the connectors that are provided with IDM 7.1. Specifically, make sure that the `connectorRef` properties reflect the new connector versions, where applicable. For example:

```
"connectorRef" : {  
  "bundleName": "org.forgerock.openicf.connectors.ldap-connector",  
  "bundleVersion": "[1.4.0.0,1.6.0.0)",  
  "connectorName": "org.identityconnectors.ldap.LdapConnector"  
},
```

Alternatively, copy the connector `.jar` files from your existing deployment into the `openidm/connectors` directory of the new installation.

Migrate UI Customizations

If you have customized the Admin UI, review any custom UI files from your IDM 7.0.x deployment (generally in the `openidm/ui/admin/extension` directory), and compare them against the corresponding IDM 7.1 files.

For each customized file, copy the corresponding default IDM 7.1 UI files to a `openidm/ui/admin/extension` directory on the new instance.

Apply your customizations to files in the new `openidm/ui/admin/extension` directory.

Chapter 5

Update the Repository

When you have migrated your configuration to the new IDM installation, you need to handle the data that is stored in your repository. There are two options to update a repository:

- Upgrade your existing IDM 7.0 repository.
- Create a new IDM 7.1 repository, then migrate your data to the new repository.

When you have upgraded the repository, or created a new repository, start the IDM server and test that all your scripts are working as expected, before migrating your data.

Upgrade an Existing Repository

Upgrading an existing repository means that you do not need to migrate data. However, you must run a series of scripts that modify the repository, to use the new features in IDM 7.1.

Note

Upgrading an existing repository is **not** supported if you use a DS repository. If you're upgrading from a previous IDM release with a DS repository, you must create a new repository, then migrate your data to the new repository.

Prepare an existing repository for IDM 7.1 as follows:

1. Shut down IDM, if it is running.
2. Clear all `configobjects` related tables. For example, in MySQL run:

```
DELETE FROM openidm.configobjects;  
DELETE FROM openidm.configobjectproperties;
```

3. From your IDM 7.1 installation, run the schema update script for your database type.

Note

You do not need to run this script if you are using a PostgreSQL repository.

The only schema update script for this release is `/path/to/openidm/db/database-type/scripts/updates/00-propertyindex.sql`. This script adds support for searchable array fields.

1. Before you run `00-propertyindex.sql`, modify the script for any properties table that is associated with a generically mapped object. Add and remove property tables, as necessary.

For each property table in the script, a new column, `propindex`, is added and the primary key is dropped and modified to include the `propindex` column.

2. When you have modified the script, run it.
3. *Optional:* If you want to use a *query filter* on any searchable array field, run the `rewriteObjects.js` script for that object path.

Important

You can only run the `rewriteObjects.js` script when the upgrade is complete, and IDM 7.1 is running, but not yet in use by any public clients.

To run the `update/rewriteObjects.js` script, specify the `objectPaths` of the arrays to make searchable. For example:

```
curl \
--header "Content-Type: application/json" \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--header 'X-OpenIDM-NoSession: true' \
--request POST \
--data-raw '{
  "type": "text/javascript",
  "file": "/path/to/openidm/bin/defaults/script/update/rewriteObjects.js",
  "globals": {
    "rewriteConfig": {
      "queryFilter": "true",
      "pageSize": 1000,
      "objectPaths": [
        "repo/config",
        "repo/internal/usermeta",
        "repo/managed/role",
        "repo/managed/user",
        "repo/reconprogressstate",
        "repo/relationships",
        "repo/scheduler/triggers"
      ]
    }
  }
}' \
"http://localhost:8080/openidm/script/?_action=eval"
```

4. Launch IDM and run the following Groovy script to clear the `reconprogressstate` data in your repository:

```
def result = openidm.query(
    "repo/reconprogresssstate", [ "_queryFilter" : "true", "_fields" : "_id" ]).result;
for ( item in result ) {
    openidm.delete("repo/reconprogresssstate/" + item["_id"], null);
}
return result.size() + " reconprogresssstate records deleted";
```

This script will work regardless of the type of repository, and can be sent as a REST call. For example:

```
curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--header "Content-Type: application/json" \
--request POST \
--data '{
  "type": "groovy",
  "source": "def result = openidm.query(\"repo/reconprogresssstate\", [ \"_queryFilter\" : \"true\",
  \"_fields\" : \"_id\" ]).result; for ( item in result ) { openidm.delete(\"repo/reconprogresssstate/\"
  + item[\"_id\"], null); }; return result.size() + \" reconprogresssstate records deleted\";"
}' \
"http://localhost:8080/openidm/script?_action=eval"
"1 reconprogresssstate records deleted"
```

5. Verify that all scripts and functions behave as expected.

Create a New Repository

Set up a new repository, following the steps in *"Select a Repository"* in the *Installation Guide*. A new repository is already configured for all the new capabilities in IDM, but does require migrating existing data to that repository.

If you create a new repository, you must still update your configuration files to use the new features.

After you have set up the new repository, migrate your data to that repository.

Chapter 6

Migrate Data

The data migration service helps you move information stored in an IDM repository to a new deployment. This service is off by default. To enable it, copy `migration.json` from `samples/example-configurations/conf/` into your `conf/` directory, and set `"enabled": true`.

Migration is run from your new installation through IDM's recon service, using your previous deployment as a data source. The data migration service supports importing information from IDM instances back to version 4. If you are migrating from a version of IDM earlier than that, you will need to follow previous update instructions to get your deployment into a state where it can be migrated using this service.

Note

Because the migration service migrates information that may be encrypted, such as passwords, you must make sure you have copied the `truststore` and `keystore` files from your previous deployment *before* you start the migration.

+ Default Data Imported by the Migration Service

- Internal Roles
- Internal Users
- Internal User Metadata
- Managed Roles
- Managed Users
- Managed Assignments
- Links and Relationships
- Scheduler jobs

Note

If you are migrating scheduler jobs from IDM 4.0 or 4.5, you will need to modify the entry in `migration.json` to be:


```
{
  "source" : "scheduler",
  "target" : "scheduler/job"
}
```

If you have additional object types (for example, managed devices), modify `migration.json` to include these objects.

- "Configure the Migration Service"
- "Run the Data Migration"

Configure the Migration Service

The data migration service is configured through `migration.json`. The default file assumes a default schema; modify the file if you have added custom managed data. The `migration.json` file can have the following properties:

enabled

Boolean, `true` or `false`. Enables the migration service.

connection

Configures the connection to the source IDM instance you are migrating from. Available properties:

instanceUrl

The URI for the source IDM instance.

authType

The authentication mechanism to the source IDM instance. Can be `basic` (username/password) or `bearer` (authentication using AM bearer tokens).

userName

Used for authenticating to the source IDM instance, if the `authType` is `basic`.

password

Used for authenticating to the source IDM instance, if the `authType` is `basic`.

clientId

Used for authenticating to the source IDM instance, if the `authType` is `bearer`.

clientSecret

Used for authenticating to the source IDM instance, if the `authType` is `bearer`.

tokenEndpoint

Used for authenticating to the source IDM instance, if the `authType` is `bearer`.

scope (optional)

List of OAuth scopes.

scopeDelimiter (optional)

Delimiter for the list of OAuth scopes.

tlsVersion (optional)

Lets you override the default TLS version.

connectionTimeout (optional)

Timeout for connecting to the source IDM instance (defaults to `10s`).

reuseConnections (optional)

Lets you override the default setting (defaults to `true`).

retryRequests (optional)

Lets you override the default setting (defaults to `true`).

hostnameVerifier (optional)

The SSL hostname verification policy. Specifies whether the host name presented by the remote server certificate is verified upon establishing new SSL connections. Possible values:

- `STRICT`: Requires that the host name match the host name presented in the certificate. Wildcards only match a single domain.
- `ALLOW_ALL`: Accepts any host name (disables host name verification).

Defaults to `STRICT`.

maxConnections (optional)

Lets you override the default maximum number of connections (default is `64`).

proxy (optional)

Lets you specify connection through a proxy server. Includes the following properties:

proxyUri

The proxy host and port to which IDM should connect.

userName

The user account to connect to the remote proxy.

password

The password of the proxy user.

socketTimeout

The TCP socket timeout, when waiting for HTTP responses. If you do not set a duration, the default is no timeout.

Example valid duration values:

- 4 days
- 59 minutes and 1 millisecond
- 1 minute and 10 seconds
- 42 millis
- unlimited
- none
- zero

mappings

A list of the endpoints that will be migrated from your old IDM instance to your new instance, expressed as mappings between the old and new instances. The complete list of mapping properties is the same as any regular [synchronization mapping](#). Properties with particular significance for data migration include the following:

source

This is the only property that is *required* for data migration. The source should be the path to the resource within the repo; for example, [repo/managed/user](#).

target

The path to the resource within the target repository. By default, this will be the same as the source path.

runTargetPhase

Specifies whether the migration should run the target phase of reconciliation. By default, this is set to `false`, as there is no data in the target repository.

sourceQuery

The query on the source system, used to find all objects to be migrated. Defaults to `"_queryFilter" : "true&fields=_id"`, which returns the IDs of all source objects.

You can improve migration performance by returning the whole source entry (setting the `sourceQuery` to `"_queryFilter" : "true"`).

sourceQueryFullEntry

(Optional). Specifies whether the defined source query returns full object data (`true`) or IDs only (`false`). Defaults to `true`.

If you do not set this parameter, IDM attempts to detect whether the full object is returned, based on the query results.

reconSourceQueryPaging

Specifies whether the migration service should use paging when querying the source IDM instance. By default, this is set to `false`. Turn paging on if you have a large data set and are concerned about memory usage.

For large data sets, you might be able to improve migration performance by turning paging on and increasing the query page size (using `reconSourceQueryPageSize`). The most effective page size will vary, depending on the available resources.

reconSourceQueryPageSize

Specifies the number of results to return per page, if paging is turned on. By default, 1000 results per page are returned.

allowEmptySourceSet

Specifies whether the migration service should continue if it encounters an empty source mapping. This is enabled by default.

properties

An array of properties you wish to perform additional actions on, such as modifying the contents of a property during the migration. (This follows the pattern you would find in a standard reconciliation. For more information about transforming data during a reconciliation, see "Transform Attributes in a Mapping" in the *Synchronization Guide*.)

policies

An array of policies you wish to apply to the data being migrated.

onCreate

The script used by the migration service for creating the data that is being migrated to the new installation. By default, this points to a Groovy script: `update/mapLegacyObject.groovy`.

onUpdate

The script used by the migration service for updating the data that is being migrated in the new installation. By default, this points to a Groovy script: `update/mapLegacyObject.groovy`.

correlationQuery

You can specify a custom correlation query. By default, this is:

```
"var map = {'_queryFilter': '_id eq \'' + source._id + '\''}; map;"
```

For more information about writing correlation queries, see "*Correlating Source Objects With Existing Target Objects*" in the *Synchronization Guide*.

validSource

You can specify a script to validate the source object prior to migration. By default, this property is empty.

endpoint

By default, the migration service endpoint is `migration`. You can use the `endpoint` property to change this if needed.

Note

Because the data migration service performs a reconciliation between your old installation and your new installation, the general reconciliation optimizations also apply to the data migration service. For more information about reconciliation optimization, see "*Tuning Reconciliation Performance*" in the *Synchronization Guide*.

Run the Data Migration

Before you run your migration, make sure that you have done the following:

- Paused any scheduled jobs on the source deployment.
- Configured your `conf/migration.json` and `update/mapLegacyObject.groovy` files on the new IDM installation.
- Moved your configuration files from the old deployment to the new one.
- If you use workflow, you must manually dump the workflow database tables, and then import them *before* you start the new instance of IDM for the first time. The workflow database tables start with

the prefix `ACT_`. For information on how to dump/import individual tables, see the documentation for your database.

When you launch the new IDM installation, a new `migration` endpoint should be available. This endpoint supports the following actions:

- `migrate`: Triggers a migration of all legacy objects from the remote system. Optionally takes a `mapping` parameter in order to specify a specific mapping to migrate. For example:

```
curl \
--header "X-OpenIDM-Username: openidm-admin" \
--header "X-OpenIDM-Password: openidm-admin" \
--header "Accept-API-Version: resource=1.0" \
--request POST \
"http://localhost:8080/openidm/migration?_action=migrate&mapping=repoManagedUser_repoManagedUser"
```

- `status`: Returns the last status for all reconciliations triggered by the migration service.
- `mappingConfigurations`: Returns the full list of migration mapping configurations.
- `mappingNames`: Returns the list of migration mapping names.

The period of time a migration takes will depend on the amount of information being migrated. Migrated data will retain the same object IDs they had in the previous deployment.

Chapter 7

Upgrade a Clustered Deployment

Follow these general steps when you are updating servers in a cluster:

- Redirect client traffic to a different IDM system or cluster.
- Shut down every node in the cluster.
- Update one node in the cluster.
- Clone the first node to the other nodes in that cluster.

IDM Glossary

correlation query	A correlation query specifies an expression that matches existing entries in a source repository to one or more entries in a target repository. A correlation query might be built with a script, but it is not the same as a correlation script. For more information, see " <i>Correlating Source Objects With Existing Target Objects</i> " in the <i>Synchronization Guide</i> .
correlation script	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 correlation query , a correlation script can be relatively complex, based on the operations of the script.
entitlement	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 assignment . A user or device with an entitlement gets access rights to specified resources. An entitlement is a property of a managed object.
JCE	Java Cryptographic Extension, which is part of the Java Cryptography Architecture, provides a framework for encryption, key generation, and digital signatures.
JSON	JavaScript Object Notation, a lightweight data interchange format based on a subset of JavaScript syntax. For more information, see the JSON site.
JSON Pointer	A JSON Pointer defines a string syntax for identifying a specific value within a JSON document. For information about JSON Pointer syntax, see the JSON Pointer RFC.

JWT	JSON Web Token. As noted in the JSON Web Token draft IETF Memo , "JSON Web Token (JWT) is a compact URL-safe means of representing claims to be transferred between two parties." For IDM, the JWT is associated with the <code>JWT_SESSION</code> authentication module.
managed object	An object that represents the identity-related data managed by IDM. Managed objects are configurable, JSON-based data structures that IDM 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 more information, see What is OSGi? Currently, only the Apache Felix container 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	IDM distinguishes between two distinct role types - provisioning roles and authorization roles. For more information, see "Managed Roles" in the <i>Object Modeling Guide</i> .
source object	In the context of reconciliation, a source object is a data object on the source system, that IDM scans before attempting to find a corresponding object on the target system. Depending on the defined mapping, IDM 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 IDM for the period during which IDM 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 IDM scans after locating its corresponding object on the source system. Depending on the defined mapping, IDM then adjusts the target object to match the corresponding source object.