

Oracle® Fusion Middleware

Integrating Enterprise Data Quality with External Systems



12c (12.2.1.4.0)

E95652-04

March 2024

The Oracle logo, consisting of the word "ORACLE" in white, uppercase, sans-serif font, centered within a solid red square.

ORACLE®

Oracle Fusion Middleware Integrating Enterprise Data Quality with External Systems, 12c (12.2.1.4.0)

E95652-04

Copyright © 2018, 2024, Oracle and/or its affiliates.

Primary Author: Oracle Corporation

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software, software documentation, data (as defined in the Federal Acquisition Regulation), or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs (including any operating system, integrated software, any programs embedded, installed, or activated on delivered hardware, and modifications of such programs) and Oracle computer documentation or other Oracle data delivered to or accessed by U.S. Government end users are "commercial computer software," "commercial computer software documentation," or "limited rights data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, reproduction, duplication, release, display, disclosure, modification, preparation of derivative works, and/or adaptation of i) Oracle programs (including any operating system, integrated software, any programs embedded, installed, or activated on delivered hardware, and modifications of such programs), ii) Oracle computer documentation and/or iii) other Oracle data, is subject to the rights and limitations specified in the license contained in the applicable contract. The terms governing the U.S. Government's use of Oracle cloud services are defined by the applicable contract for such services. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle®, Java, MySQL, and NetSuite are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Inside are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Epyc, and the AMD logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information about content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc>.

Contents

Preface

Audience	viii
Documentation Accessibility	viii
Related Documents	viii
Conventions	ix

1 Integrating with Subversion

1.1 Software Requirements	1-1
1.2 Understanding the Integration Architecture	1-1
1.3 Setting Up a Repository	1-4
1.4 Configuring EDQ with Subversion	1-4
1.4.1 Configuring a New EDQ Installation	1-5
1.4.2 Retaining Existing Configuration Information	1-5
1.5 Understanding the Integration Elements	1-6
1.6 Reviewing a Deployment Example	1-7
1.7 Troubleshooting Errors	1-9

2 Integrating with Git

2.1 Understanding the Integration Architecture	2-1
2.2 Preparing the Git Workspace	2-2
2.3 Configuring EDQ with Git	2-3
2.4 Using EDQ	2-3

3 Integrating with IBM Global Name Recognition

3.1 System Requirements	3-1
3.2 Configuring the EDQ Server	3-1
3.3 Building the Search Library	3-2
3.4 Configuring the GNR Connector	3-3
3.4.1 Creating the EDQ GNR Properties File	3-3
3.5 Creating the Search Configuration Files	3-4

4 Integrating with Experian QAS

4.1	Software Requirements	4-1
4.2	Integrating with Experian QAS	4-1
4.3	Migrating QAS integrations	4-2

5 Integrating with Capscan Matchcode

5.1	Software Requirements	5-1
5.2	Integrating the Capscan Matchcode Libraries into EDQ	5-1
5.3	Customizing the Matchcode API	5-2

6 Using the Command Line Interface

6.1	Running the Command Line Interface	6-1
6.2	Understanding the Commands and Arguments	6-1
6.2.1	runjob	6-1
6.2.2	runopsjob	6-2
6.2.3	droporphans	6-3
6.2.4	listorphans	6-3
6.2.5	scriptorphans	6-3
6.2.6	list	6-3
6.2.7	showlogs	6-3
6.2.8	shutdown	6-4
6.2.9	version	6-4
6.3	Reviewing Examples	6-4
6.3.1	Listing All the Available Commands	6-4
6.3.2	Listing the Available Parameters for a Command	6-5
6.3.3	Running a Named Job	6-5
6.3.4	Running a Named Job in Operations Mode	6-5

7 Configuring Additional Database Connections

7.1	Using JNDI to Connect to Data Stores	7-1
7.2	Connecting to an Oracle Database Using tnsnames.ora	7-1
7.2.1	To Configure EDQ to Connect Through TNS	7-1
7.3	Connecting to an Oracle Database Using Oracle Internet Directory (LDAP)	7-2

8 Configuring EDQ to Process XML Data Files

8.1	Using Simple XML Data Stores	8-1
8.1.1	Reading Simple XML Files	8-1
8.1.2	Writing Simple XML Files	8-2
8.2	Using XML and Stylesheet Data Stores	8-2
8.2.1	Using DN-XML	8-3
8.2.2	Reading Custom XML Files	8-4
8.2.2.1	Configuring the Data Store	8-6
8.2.3	Writing Custom XML Files	8-6
8.2.3.1	Configuring the Data Store	8-7

9 Using the EDQ Configuration API

9.1	REST Interface for Projects	9-1
9.1.1	Retrieving a List of EDQ Projects	9-1
9.1.2	Creating a Project	9-2
9.1.3	Deleting a Project	9-2
9.2	REST Interface for Data Stores	9-3
9.2.1	Retrieving a List of Data Stores	9-3
9.2.2	Creating a Data Store	9-4
9.2.3	Deleting a data store	9-6
9.3	REST Interface for Snapshots	9-6
9.3.1	Retrieving a List of Snapshots	9-7
9.3.2	Creating a Snapshot	9-7
9.3.3	Deleting a Snapshot	9-8
9.4	REST Interface for Processes	9-8
9.4.1	Retrieving a List of Processes	9-9
9.4.2	Deleting a Process	9-9
9.4.3	Creating a Simple Process	9-10
9.5	REST Interfaces for Jobs	9-11
9.5.1	Retrieving a List of Jobs	9-11
9.5.2	Deleting a Job	9-12
9.5.3	Creating a Simple Job	9-12
9.5.4	Running a Job	9-12
9.5.5	Cancelling a Running Job	9-13
9.5.6	Getting the Status of a Job	9-14
9.5.7	Getting the Details of All Running Jobs	9-14
9.6	REST Interface for Reference Data	9-15
9.6.1	Retrieving a List of Reference Data	9-15
9.6.2	Retrieving Contents of Reference Data	9-16
9.6.3	Creating Reference Data	9-17

9.6.4	Deleting Reference Data	9-17
9.7	REST Interface for Web Services	9-18
9.7.1	Retrieving a List of Web Services	9-18
9.7.2	Creating or Updating a Web Service	9-19
9.7.3	Deleting a Web Service	9-20
9.8	Example: Profiling from an External Application	9-20

10 Using REST APIs for Importing and Exporting Configuration Objects

10.1	REST Interface for Packaging Configuration	10-1
10.2	JSON Payload Format	10-2
10.3	Packaging Task Status Result Format	10-6
10.4	Packaging REST API Triggers	10-8

11 Using REST APIs for User Management and Launchpad Administration

11.1	REST Interface for Creating and Updating Groups	11-1
11.2	REST Interface for External Group Mappings	11-3
11.3	REST Interface for Creating and Updating Users	11-4
11.4	REST Interface for Launchpad Applications	11-6

12 Using the Java Messaging Service (JMS) with EDQ

12.1	Understanding the Message Queue Architecture	12-1
12.2	Uses of JMS with EDQ	12-1
12.3	Configuring EDQ to Read and Write JMS Messages	12-2
12.4	Defining the Interface Files	12-2
12.4.1	Knowing the <attributes> section	12-2
12.4.2	Knowing the <messengerconfig> Section	12-3
12.4.3	Knowing the <messagebody> section	12-5
12.5	Illustrations	12-5

13 Using Apache Kafka with EDQ

13.1	Introduction to Kafka and EDQ	13-1
13.2	Configuring EDQ to Read and Write Kafka Records	13-1
13.3	Defining the Interface Files	13-2
13.3.1	Understanding the <attributes> section	13-2
13.3.2	Understanding the <messengerconfig> section	13-3
13.3.3	Understanding the <incoming> or <outgoing> section	13-3
13.3.3.1	Understanding the <messageheaders> section	13-4
13.3.3.2	Understanding the <messagebody> section	13-4

13.4	Illustrations	13-6
------	---------------	------

14 Using Amazon Simple Queue Service (Amazon SQS) with EDQ

14.1	Introduction to Amazon SQS and EDQ	14-1
14.2	Configuring EDQ to Read and Write Amazon SQS Messages	14-1
14.3	Defining the Interface Files	14-2
14.3.1	Understanding the <attributes> section	14-2
14.3.2	Understanding the <messengerconfig> section	14-3
14.3.3	Understanding the <incoming> or <outgoing> section	14-3
14.3.3.1	Understanding the <messageheaders> section	14-4
14.3.3.2	Understanding the <messagebody> section	14-5
14.4	Illustrations	14-5

15 Using Oracle Cloud Infrastructure (OCI) Queue with EDQ

15.1	Introduction to OCI Queue and EDQ	15-1
15.2	Configuring EDQ to Read and Write OCI Queue Messages	15-1
15.3	Defining the Interface Files	15-2
15.3.1	Understanding the <attributes> section	15-2
15.3.2	Understanding the <messengerconfig> section	15-3
15.3.3	Understanding the <incoming> or <outgoing> section	15-3
15.3.3.1	Understanding the <messageheaders> section	15-3
15.3.3.2	Understanding the <messagebody> section	15-4
15.3.4	Illustrations	15-4

16 Using Scripted Global Web Services with EDQ

16.1	Introduction to Global Web Services and EDQ	16-1
16.2	Configuring EDQ to Read and Write Web Service Requests	16-1
16.3	Defining the Interface Files	16-2
16.3.1	Understanding the <attributes> section	16-2
16.3.2	Understanding the <messengerconfig> section	16-3
16.3.3	Understanding the <incoming> or <outgoing> section	16-3
16.3.3.1	Understanding the <messageheaders> section	16-3
16.3.3.2	Understanding the <messagebody> section	16-4
16.4	Illustrations	16-5

Preface

Describes how to integrate Enterprise Data Quality with external systems and applications.

Audience

This document is intended for advanced users of EDQ and administrators responsible for integrating EDQ with third-party applications.

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc>.

Accessible Access to Oracle Support

Oracle customers who have purchased support have access to electronic support through My Oracle Support. For information, visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> or visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> if you are hearing impaired.

Related Documents

For more information, see the Oracle Enterprise Data Quality documentation set.

Find the latest version of the EDQ guides and all of the Oracle product documentation at

<https://docs.oracle.com>

Online Help

Online help is provided for all EDQ user applications. It is accessed in each application by pressing the **F1** key or by clicking the Help icons. The main nodes in the Director project browser have integrated links to help pages. To access them, either select a node and then press **F1**, or right-click on an object in the Project Browser and then select **Help**. The EDQ processors in the Director Tool Palette have integrated help topics, as well. To access them, right-click on a processor on the canvas and then select **Processor Help**, or left-click on a processor on the canvas or tool palette and then press **F1**.

Conventions

The following text conventions are used in this document:

Convention	Meaning
boldface	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.
<i>italic</i>	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
<code>monospace</code>	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.

1

Integrating with Subversion

This chapter describes how to integrate and use EDQ with the Subversion version control system.

The following sections are included:

- [Software Requirements](#)
- [Understanding the Integration Architecture](#)
- [Setting Up a Repository](#)
- [Configuring EDQ with Subversion](#)
- [Understanding the Integration Elements](#)
- [Reviewing a Deployment Example](#)
- [Troubleshooting Errors](#)

1.1 Software Requirements

EDQ supports integration with all current versions of Subversion. For more information about Subversion, see the Apache Subversion website found at <http://subversion.apache.org/>.

The Subversion server with which EDQ is being integrated must meet these prerequisites:

- Support Hypertext Transfer Protocol (HTTP) and Distributed Authoring and Versioning (DAV) access.
- Require authentication on commit.
- Not require authentication on checkout or update.

When Subversion is integrated with EDQ as a store of configuration information, the following restrictions and limitations apply. Consider the following points before deciding to configure integrated version control using Subversion.

- You cannot update or revert an item that is open in the Director interface or the Subversion server.
- You cannot rename a project once the project is under version control. This is critical in avoiding duplication of reference processor names in a project.
- Deleting a project does not remove it from the Subversion repository.
- Case insensitive name matching is used.

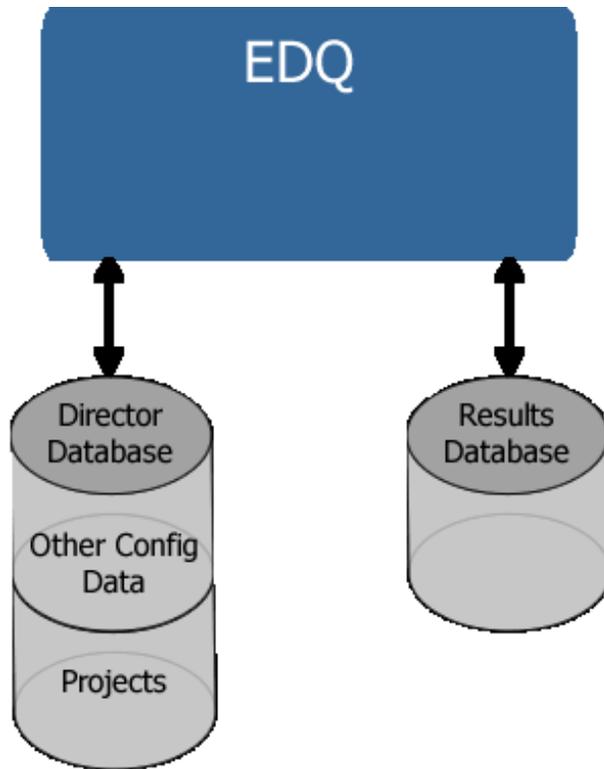
1.2 Understanding the Integration Architecture

The EDQ server can be configured to be aware of a Subversion server as a store of configuration information.

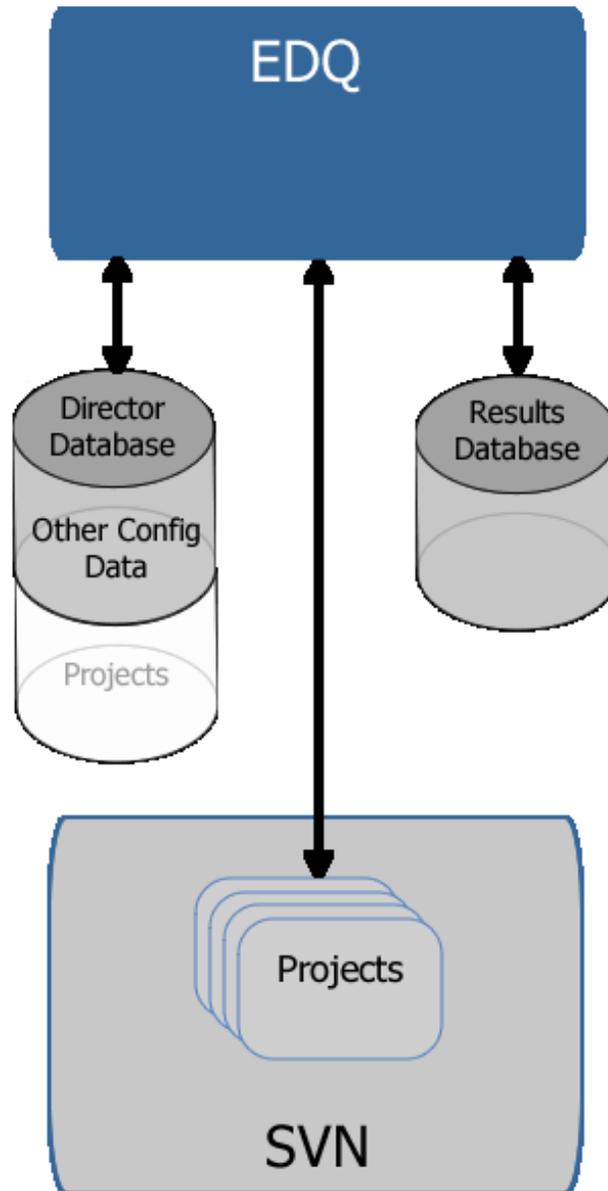
 **Note:**

In this instance, configuration information means information that is managed using the Director UI; for example, projects and system-level data.

In a standard EDQ instance, configuration information, including project information, is stored in the Director database:



The following figure shows an EDQ instance integrated with Subversion:



Note:

The Director database is still required because it contains data derived from the file-mastered configuration that has been normalized to allow querying by the applications.

With EDQ configuration files mastered and stored in a Subversion repository, a Subversion client can be used to commit or otherwise access them. Because EDQ includes an embedded Subversion client, Subversion client operations to control configuration changes

can be performed directly in Director once the EDQ integration with Subversion has been enabled.

1.3 Setting Up a Repository

The first stage of configuration is to create a workspace directory where the checked out data is stored:

1. Create a directory on the disk where desired (for example, `C:\MyRepository`) and then add it and commit it to Subversion.
2. Inside the newly created directory, set the following Subversion property:

```
svn propset edq:systemversion 12.1.3:base .
```

 **Note:**

Set Subversion property to "12.1.3:base" and not to the current version of EDQ.

3. Commit these changes into Subversion. Your workspace now displays these properties:

```
svn proplist -v .  
Properties on '.':  
  edq:systemversion  
  12.1.3:base
```

4. Create the following subdirectories in the newly created directory:
 - Data Stores
 - Hidden Reference Data
 - Images
 - Projects
 - Published Processors
 - Reference Data
5. Add and commit these directories. The repository is now set up correctly for EDQ.

The preceding steps only need to be performed once per repository. All remaining changes can be made using EDQ.

1.4 Configuring EDQ with Subversion

Subversion must be integrated with a fresh installation of EDQ.

 **Caution:**

When an EDQ instance is integrated with Subversion, all pre-existing and other configuration information is lost. To retain this information, you must package and export it first. For further details, see [Retaining Existing Configuration Information](#).

 **Note:**

Oracle recommends that a single workspace be assigned to each instance of EDQ because it is difficult to move between workspaces in a single EDQ instance.

1.4.1 Configuring a New EDQ Installation

To configure a new EDQ installation:

1. Shut down the application server.
2. Check out the workspace from Subversion. It is not necessary to checkout the whole tree; just the workspace directory itself is required.
3. Edit the `director.properties` file in the `ORACLE_HOME/user_projects\domains\domains\edq_domain\edq\oedq.local.home` directory.
4. Add the following line replacing the directory path with that of the absolute path to your root workspace directory. For example:

```
scs.workspace = C:\MyRepository
```

 **Note:**

This example demonstrates the need to escape colon (:) and backslash (\) characters in the path with a backslash. You must also escape space characters in the path with a backslash.

5. Start your EDQ server, and then start Director.
6. Check the top of the `Main0.log` file for an `INFO` message listing the name of the SCCS workspace you added. For example:

```
INFO: 02-Sep-2013 10:05:21: SCCS workspace is C:\MyRepository
```
7. If no errors follow this message, EDQ is configured to use Subversion. If there are errors, see [Troubleshooting Errors](#), for possible solutions.

1.4.2 Retaining Existing Configuration Information

As previously stated, Subversion must be integrated with a fresh installation of EDQ. Therefore, any pre-existing projects and other configuration items in an EDQ installation must be packaged before integration begins and then imported to the new installation afterwards:

1. Package all configuration items in the current EDQ instance into DXI files.

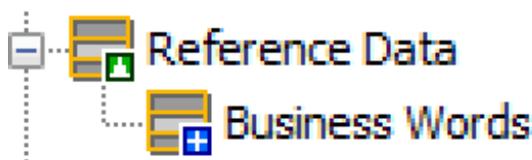
2. Install a new instance of EDQ with the Subversion integration enabled.
3. Import the DXI files into the new instance, and commit the files to the Subversion workspace.
4. Check that the configuration items are all valid and working correctly.
Note that all passwords for Data Stores must be re-entered after a configuration import.
5. Decommission the previous instance.

1.5 Understanding the Integration Elements

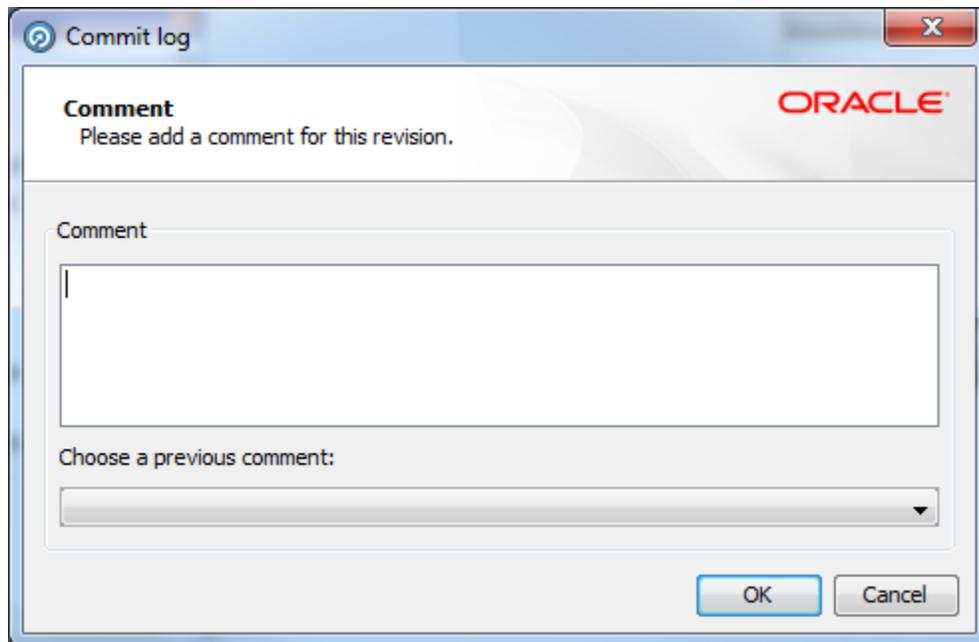
Once EDQ is integrated with Subversion enabled, the following interface elements become visible within the Director application:

- Subversion status icon overlays in Project Browser - There are two icons used to indicate the three possible Subversion statuses of nodes in the Project Browser:
 - No icon - The node (and its sub-nodes) are all up to date.
 - Green icon- This node (and its sub-nodes) have modifications.
 - Blue icon - This node (and its sub-nodes) is new/currently not under Version Control.

For example, the following image shows both icons in use. The Reference Data node is modified (green icon) as one of its sub-nodes has changed. A new piece of Reference Data, Business Words, has been added, and is marked with the blue icon:

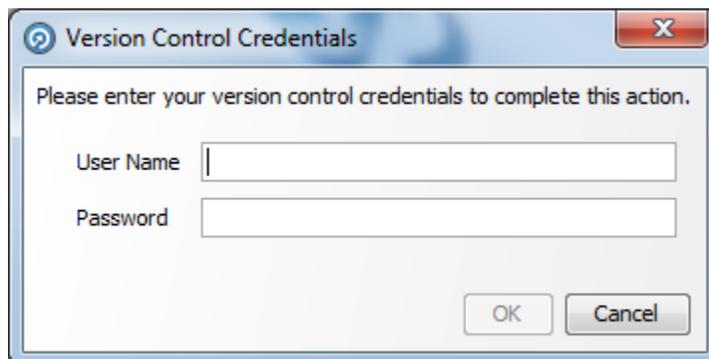


- Version Control tab - The Properties dialog (displayed by right-clicking on an item in the Project Browser and selecting **Properties**) now contains a Version Control tab that describes the state of the item, when it was last updated, its Subversion revision, and whether it is current.
- New context menu for Version Control - The Project Browser right-click menu now contains a Version Control option. When selected, this displays a sub-menu with Subversion options to update, commit, revert, compare or view the log for the item. These options are recursive. For example, if you perform View Log on a single process then you will see the log for this process only, but if you perform View Log on the Processes node you will see changes for all processes.
- Comment and credentials dialogs on commit - When you commit changes to the repository, Director displays the Commit log dialog:



In this dialog you can enter a comment describing the change in the Comment field. Alternatively, you can automatically populate the field by choosing a comment from the list of comments previously entered in the current session.

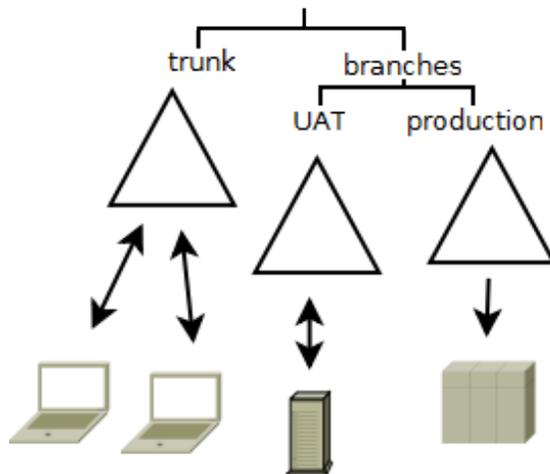
After you click **OK** in the Commit log dialog, Director displays the Version Control Credentials dialog if you have not already provided your credentials in the current session:



In this dialog you enter your user name and password for the Subversion repository and then click **OK**.

1.6 Reviewing a Deployment Example

An example deployment is presented here. In this illustration, there is a single Subversion server that holds three copies of the configuration for four EDQ installations:



The copies of the configuration are:

- **trunk** - the traditional location that all development work is performed on. New features of the configuration are developed and saved here.
- **branches and UAT** - this branch represents the copy of the configuration under UAT testing.
- **branches and production** - this branch represents the production copy of the configuration.

The four EDQ installations using the Subversion server for storing their configuration are:

- Two development laptops where design work and maintenance of existing projects are carried out.
- A UAT server for User Acceptance Testing changes.
- A production server for production runs.

In this example deployment, the laptop users develop configuration for individual projects on their own laptops and then commit changes back to the subversion repository on "trunk". Where the developers are co-operating on developing a project they will periodically update their local installation to pick up changes from the other developer.

At some point development reaches a point where it needs to be released to UAT for testing. A release manager then copies the necessary projects from "trunk" to "UAT" on the subversion server.

For example, the following Subversion command may be used:

```
svn cp -m"Release Project X to UAT" http://svn/repos/config/trunk/ProjectX  
http://svn/repos/config/branches/UAT
```

The test manager then updates the UAT server's projects to load the new configuration into the EDQ server. Over a period of time testing continues. As issues are found they are fixed in the UAT environment and committed back to the subversion repository.

Once UAT environment has achieved an acceptable test level it is promoted to release. This achieved in much the same way as the release from development to UAT. The necessary projects are copied across in the version control repository and then the production server is updated to use this configuration.

1.7 Troubleshooting Errors

You may encounter the following errors for which the cause and solution is provided.

Error	Cause and Solution
Configuration database is not compatible with workspace	The database has been used with a different workspace. This error usually arises occurs when operations have been performed in EDQ before Subversion version control is enabled. There are two solutions: drop and recreate the Director database or reinstall EDQ.
Unable to open an ra_local session to URL	This may occur when trying to commit files to an invalid repository. The EDQ integration is not compatible with file-based repositories (those repositories beginning with file:/// or C:\example). A fully declared http:// path to the repository must be made.

2

Integrating with Git

This chapter describes how to integrate and use EDQ with the Git version control system. Consider the following points before deciding to configure integrated version control using Git.

- Project and global objects are stored in the file systems, but Software Configuration Management (SCM) operations such as commits, pushes and pull requests are performed manually outside of EDQ.
- The EDQ system must be freshly initialized. The Git configuration must be completed before you start the system for the first time.

Note:

If you have an existing EDQ system with configuration that you want to convert to using Git, you should package all projects on this system and back up its local home. You can then import the projects and any configuration extensions or other needed files in the local home onto a Git-integrated system.

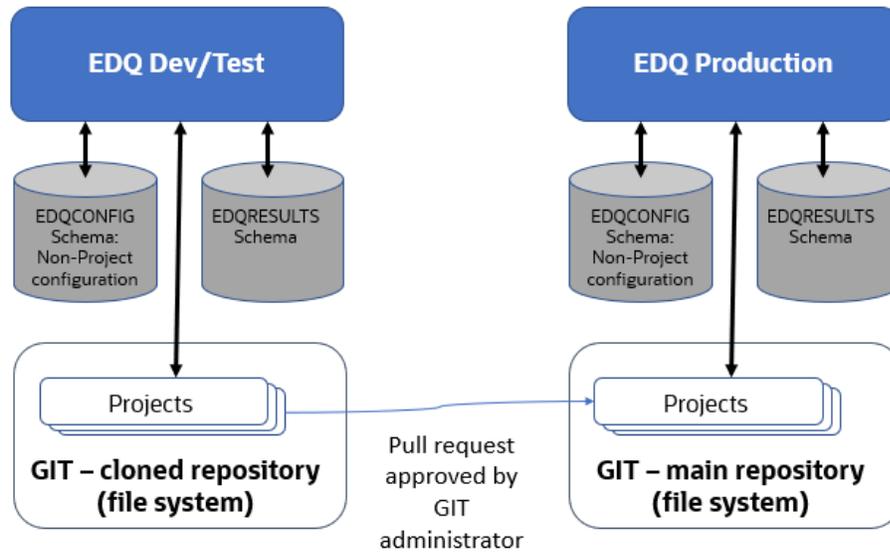
The following sections are included:

- [Understanding the Integration Architecture](#)
- [Preparing the Git Workspace](#)
- [Configuring EDQ with Git](#)
- [Using EDQ](#)

2.1 Understanding the Integration Architecture

The EDQ server can be configured to be aware of a Git server as a store of configuration information.

The following figure illustrates a typical setup of two EDQ instances that are integrated with Git. Here the Dev/Test instance is used to make and test changes. The tested changes are then promoted onto a main branch used by a Production instance by a pull request that is approved by a Git administrator. Note that the Dev and Test instances could equally be separate. For example, they may both work on the same fork, but with changes only moving to Test once committed and pushed from Dev.



2.2 Preparing the Git Workspace

The first stage of configuration is to create a workspace directory where the checked out data is stored:

1. Create a new empty directory in a Git workspace and create a `.wsprops` file that includes this single line:

```
systemversion=12.1.3:base
```

Note:

Set the property to "12.1.3:base" and not to the current version of EDQ. This is the correct value for all versions of EDQ 12.2.x.

2. Create the following subdirectories in the newly created directory:
 - Data Stores
 - Hidden Reference Data
 - Images
 - Projects
 - Published Processors
 - Reference Data
3. In each directory, create an empty `.gitignore` file to ensure that the directories can be committed to Git successfully. You need to do this because you cannot commit empty directories to Git. To create and populate the directories, use the following script:

```
rootdirs="Data_Stores Hidden_Reference_Data Images Projects
Published_Processors
Reference_Data"
```

```
# Create root dirs
for i in $rootdirs
do x=$(echo $i | tr '_' ' ')
  mkdir "$x"
  echo > "$x/.gitignore"
  echo created $x
done
cat > .wsprops <<EOF
systemversion=12.1.3:base
EOF
```

4. Add and commit these directories to Git. The repository is now set up correctly for EDQ.

2.3 Configuring EDQ with Git

Git must be integrated with a fresh installation of EDQ. Before you start the EDQ server for the first time, edit `director.properties` and add these lines:

```
sccs.workspace = file system path to root directory
sccs.vcs.type = null
```

Here, the root directory in the first line is the directory where you created the repository in [Preparing the Git Workspace](#). The second line disables the integrated Subversion support.

After you have edited `director.properties`, start the EDQ server.

The built-in reference data objects will be created in the root Reference Data folder. Commit these objects to Git and push them to the origin server.

2.4 Using EDQ

After the Git configuration is complete, you can use the EDQ Director client to create and edit global and project objects as normal. The objects are stored in the file system location set by the `sccs.workspace` property in [Configuring EDQ with Git](#).

To make changes visible to other users of the Git repository, commit and push the changes using the standard Git command line tools. For example:

```
$ cd /opt/git/repo/dev
$ git add Projects/test1
$ git commit -m 'Committing a project'
$ git push
```

After the changes have been pushed to the origin server, other users can clone or pull from the repository to work on the same objects.

To update the local workspace with changes created by other users, use a `git pull` command:

```
git pull
remote: Counting objects: 17, done.
remote: Compressing objects: 100% (8/8), done.
remote: Total 13 (delta 3), reused 0 (delta 0)
```

```
Unpacking objects: 100% (13/13), done.  
...
```

If the local EDQ server is running, rescan the workspace for changes. To do this, use the `jshell.jar` utility to run the following script:

```
$ java -jar jshell.jar scripts/sccs/scan.groovy \  
      -server host:8090 -user username -pw password
```

The Project Browser of running EDQ clients reflect the changes that are picked up by the script.

Note that the `scan.groovy` script may update process and job objects that are open in EDQ clients. However, the Canvas does not reflect these changes. You must close and reload the objects to see the changes. A good practice is to ensure that no clients or jobs are running when a scan is run. After you execute any Git operation that changes the workspace contents, such as switching branches, make sure that you run the `scan.groovy` script or restart the EDQ server.

3

Integrating with IBM Global Name Recognition

This chapter describes how to integrate with IBM Global Name Recognition (GNR). You can configure EDQ to connect to IBM GNR to facilitate linguistic analysis of names, and linguistically sensitive name searching.

This chapter includes the following sections:

- [System Requirements](#)
- [Configuring the EDQ Server](#)
- [Building the Search Library](#)
- [Configuring the GNR Connector](#)
- [Creating the Search Configuration Files](#)

3.1 System Requirements

To enable EDQ connectivity with IBM GNR, you must have the following:

- EDQ 12c (12.2.1.1.0) installed on 64-bit AIX operating system or Linux operating system running 64-bit Java.
- IBM GNR 4.2.1 (4.2 + 4.2.1 fixpack) or later, including the hotfix based on GNR 4.2.2 (4.2 +4.2.2 fixpack). For more information, see the IBM website at <http://www.ibm.com>.

EDQ does not make use of any of the web services provided by GNR, so you do not need to configure these during GNR installation.

Note:

GNR can only be installed on an EDQ instance if you have the required license agreements with both Oracle and IBM.

3.2 Configuring the EDQ Server

The `LD_LIBRARY_PATH` must be set as required for the installation environment.

The EDQ GNR analytic processors use a shared library (`.so`) in the `lib64` directory of the GNR installation. This directory must be specified in an environment variable passed to the EDQ server.

In a Linux 64-bit environment, the environment variable is `LD_LIBRARY_PATH`; for example:

```
LD_LIBRARY_PATH=$LD_LIBRARY_PATH:gnr-installation-dir/lib64
```

`gnr-installation-dir` is the GNR installation path; for example, `/opt/GNR/GNM`.

In an AIX environment, the environment variable name is `LIBPATH` instead of `LD_LIBRARY_PATH`.



Note:

The environment variable must be available to the application server process.

3.3 Building the Search Library

The GNR Search Processor uses a native library that must be linked with the GNR libraries.

Oracle supplies these files to create the library:

- Two `Makefile` templates, one for each platform, that script the building of the search library
- The `namehunter.o` object module file

When building the library on an AIX system, IBM C++ must be available. When building the library on Linux systems, the GCC C++ compiler must be available and it must be of the same version used to create the GNR libraries, as specified in the IBM GNR documentation.

The `Makefile` template for a 64-bit AIX operating system is as follows:

```
# Build library from object file
# -----

GNR=/opt/GNR/GNM
CFLAGS=-qmkshrobj
LIBS=-lNameHunter -lNameTransliterator -lsicu18n -lsicuuc -lsicudata
SDK=aix61_64-xlc9-release
LIBDIRS=-L$(GNR)/sdk/$(SDK)/lib -L$(GNR)/sdk/icu4c/$(SDK)/lib

all: libnimrod.so

libnimrod.so: namehunter.o
        xlc++_r -q64 $(CFLAGS) -o $@ $? $(LIBDIRS) -lNameHunter -
lNameTransliterator
-lsicu18n -lsicuuc -lsicudata
```

The `Makefile` template for 64-bit Linux operating system is as follows:

```
# Build 64-bit library from object file
# -----

GNR=/opt/GNR/GNM
CFLAGS=-shared -fPIC

all: libnimrod.so

libnimrod.so: namehunter.o
        g++ -m64 $(CFLAGS) -o $@ $? -L$(GNR)/sdk/rhel4_64-gcc34-release/lib
-L$(GNR)/sdk/icu4c/rhel4_64-gcc34-release/lib -lNameHunter -lNameTransliterator
-lsicu18n -lsicuuc -lsicudata
```

Before running the `Makefile` script for your platform, ensure that the value of `GNR` in the `Makefile` template is set correctly, according to the GNR installation directory location. On the AIX operating system, also ensure that the value of `SDK` is set correctly, according to the system architecture.

After running the edited `Makefile`, the newly created `libnimrod.so` shared library file can be installed anywhere and can be copied to other systems with GNR installs.

3.4 Configuring the GNR Connector

The EDQ GNR connector requires three types of configuration files to integrate it with GNR:

- The `gnr.properties` properties file in the EDQ installation
- The `nameworks.config` configuration file in the GNR installation
- Search configuration files in the EDQ installation

3.4.1 Creating the EDQ GNR Properties File

The `gnr.properties` file and the `gnr` subdirectory that contains it must be manually created and placed in the `gnr` subdirectory of the EDQ configuration directory. It must contain the following properties:

`gnr.install`

The GNR installation path. This is the path to the directory containing the following GNR subdirectories:

- `bin`
- `bin64`
- `data` (which contains the GNR data files)
- `lib`
- `lib64`

`analytics.config`

The absolute location of the `nameworks.config` configuration file in the GNR installation.

`search.jnilib`

The absolute location of the `libnimrod.so` shared library, which was built using the `Makefile` template.

`nameworks.config`

During GNR installation, a `nameworks.config` file is created and stored in the GNR data directory.

The critical part of the `nameworks.config` file is the reference files section:

```
[Reference Files]
NameSifter=/opt/GNR/GNM/data/SifterRules.ibm
```

The `NameSifter` value must refer to the `SifterRules.ibm` file in the GNR installation.

3.5 Creating the Search Configuration Files

Search configuration files are located in the `gnr/search` subdirectory of the EDQ configuration directory. They are read by the connector and used to set parameters for the Search function.

A sample search configuration file named `search.config` is available in the `support/data/search` subdirectory of the EDQ installation. To create a search configuration file, copy this sample file to the `gnr/search` subdirectory of the EDQ configuration directory and edit the copy to suit your needs.

3.5.1 Support for GNR 3.2 and GNR 4.2 in Search Configuration Files

The search configuration format changed slightly from GNR 3.2 to GNR 4.2, and the EDQ GNR connector supports both versions as far as possible. It also processes data for Organization searches.

The basic differences between the search configuration files in GNR 3.2 and GNR 4.2 are:

- GNR 4.2 specifies the parameter files (for example, tags and variants) in the `[hunter]` section. GNR 3.2 uses the `[search]` section. The EDQ GNR connector looks in the `[hunter]` section first then the `[search]` section.
- The tag and variant files in GNR 4.2 are specified by keys such as `ibmTaqFile` and `custTaqFile`. In GNR 3.2, just `taqFile` is used. The EDQEDQ GNR connector looks for `taqFile`, `ibmTaqFile` and `custTaqFile` and loads each if found. The same rules are used for variant and terms files.
- The generic reg file is set by a specific `genericRegFile` setting; in GNR 3.2 this always defaults to the `anglo` reg file.
- Some settings have been added to the `[parms]` sections, and others have been removed.

4

Integrating with Experian QAS

This chapter describes how to integrate with Experian QAS API and migrate earlier versions of the QAS integration to EDQ versions 8.1.3 or later.

This chapter includes the following sections:

- [Integrating with Experian QAS](#)
- [Migrating QAS integrations](#)

4.1 Software Requirements

EDQ includes a connector to the Experian QAS Batch API. You must have an installed version of the Experian QAS Batch API software appropriate for your platform. For more information, see the Experian Data Quality website found at <http://www.qas.com/>.

4.2 Integrating with Experian QAS

Integration of EDQ with Experian QAS Batch is carried out by editing the `qas.properties` file distributed with EDQ. The `qas.properties` file is located in the `ORACLE_HOME/user_projects/domains/domains\edq_domain\edq\oedq.home\qas` directory.

Once both EDQ and the QAS Batch software have been installed, edit the properties in the `qas.properties` as required for your integration. The properties specified in the file are as follows:

Property name	Description	Default Value
<code>qas.install.path</code>	The location of the Experian QAS Batch installation.	C:\\Program Files (x86)\\QAS\\QuickAddress Batch API
<code>qas.qaworld.ini</code>	The path to the <code>QAWorld.ini</code> file to use. This allows you to create and edit copies of the original <code>QAWorld.ini</code> file distributed with Experian QAS Batch API. If no value is specified for this property, it uses the <code>QAWorld.ini</code> file within the Experian QAS Batch API installation.	None

Property name	Description	Default Value
<code>max.number.connections</code>	The maximum number of connections that EDQ will create to the API. This should not be set to a value greater than 32. Note: There is a known threading and memory issue with Experian QAS Batch API versions 6.85, 6.89 and 6.95, where if maximum usage of 32 instances and 8 threads is reached, the Batch API may crash. This can be avoided by setting this property to a value between 18 and 22 inclusive.	32
<code>connection.pool.timeout</code>	This property specifies the number of milliseconds that a connection can be idle for before it is closed by the pool management functionality. If this is set to -1, idle connections will not be closed.	60,000
<code>connection.pool.timer.interval</code>	This property specifies how often, in milliseconds, the connection pool is scanned for idle connections.	60,000
<code>default.layout</code>	The default layout to use. If no layout with the specified name is available in the specified <code>QAWorld.ini</code> file, this property is ignored.	GBR

4.3 Migrating QAS integrations

Some earlier versions of EDQ (versions prior to 8.1.3) were shipped with a customized version of `QAWorld.ini` that was used instead of the version contained within Experian QAS. It is not possible to migrate automatically from these versions of the Experian QAS integration to the later versions. To migrate an earlier Experian QAS integration, you must:

- Locate the local copy of `QAWorld.ini` and copy any custom layouts specified in this file into the version of `QAWorld.ini` contained within Experian QAS.
- Update the settings in `qas.properties`. See [Integrating with Experian QAS](#).

In addition, any existing processes and results books that make use of QAS processors must be updated as follows:

- Open each configured QAS processor and rename the output attributes to match those in the new `QAWorld.ini` file.
- Open any results books built on results grids from QAS processors and re-map the fields to the new output attribute names.

Refer to the following table for the affected releases for each version of EDQ. All versions of EDQ (previously known as dn:Director) prior to 7.2 are affected:

Version	Releases affected
EDQ (dn:Director) 7.2	Release 7.2.9 and all earlier 7.2 releases

Version	Releases affected
EDQ (dn:Director) 8.0	Release 8.0.21 and all earlier 8.0 releases
EDQ (dn:Director) 8.1	Release 8.1.2 and all earlier 8.1 releases

5

Integrating with Capscan Matchcode

This document describes how to integrate address verification and cleansing features from GBGroup Capscan Matchcode with EDQ. This documentation is intended for system administrators responsible for installing and maintaining EDQ applications. This chapter includes the following sections:

- [Software Requirements](#)
- [Integrating the Capscan Matchcode Libraries into EDQ](#)
- [Customizing the Matchcode API](#)

5.1 Software Requirements

You must have the Capscan Matchcode software installed on a system that is accessible to the EDQ Server. For more information, see the GBGroup Matchcode website found at <https://www.gbgroup.com/products/>.

5.2 Integrating the Capscan Matchcode Libraries into EDQ

EDQ includes a connector to the Capscan Matchcode API. EDQ includes a connector to the Capscan Matchcode API from GBGroup. This API provides address verification and cleansing features. This API provides address verification and cleansing features. Integrate the Capscan Matchcode Libraries into EDQ as follows:

1. Copy the `capscan.jar` client API file from the Capscan Matchcode installation to the `ORACLE_HOME/user_projects/domains/edq_domain/servers/edq_server1/tmp/_WL_user/edq/iz3lfy/war/WEB-INF/widgetjars` directory.

The location of the `capscan.jar` file in a Capscan Matchcode installation depends on the installation platform; for example, it is located in the `Capscan\SDK\Matchcode client API\Java` directory on Windows.

Note:

You can copy the `capscan.jar` file to a directory other than the default directory given in this step. If you do so, you must edit the `capscan.jar` property of the `capscan.properties` file to specify the location of the file. If you specify a relative path, the path must be relative to one of the directories in the EDQ configuration path.

2. Edit `capscan.properties` file in the `ORACLE_HOME/user_projects/domains/edq_domain/edq/oedq.home/capscan` directory.
3. Edit the `server.host` property to refer to the system where Capscan Matchcode is running.

4. Restart your Server.

5.3 Customizing the Matchcode API

Various aspects of the Capscan Matchcode API behavior can be controlled using the `capscan.properties` file. It allows you to set the following properties:

capscan.jar

The location of the `capscan.jar` client API file in the installation. If you specify a relative path, the path must be relative to one of the directories in the configuration path. The default value is the relative path `capscan/capscan.jar`.

server.host

The IP address of the machine where Capscan Matchcode is running.

connection.timeout

A timeout period, in seconds, after which the Capscan Matchcode API will abort the search and return (a timeout period of zero indicates that there is no time limit on searches).

connection.type

The connection mode to use when communicating with the Capscan Matchcode API. The default connection mode is `CONNECTIONLESS`. The remaining options are:

- `CONNORIENTED`
- `STATELESS`
- `WEBCONNECTION`
- `ONDEMAND`

For information about these connection modes, refer to the Capscan Matchcode API documentation.

number.capscan.connections

The number of connections that EDQ should make to the Capscan Matchcode API.

number.threads

The number of threads that should be used when communicating with the Capscan Matchcode API.

The default contents of the `capscan.properties` file are as follows:

```
# This configuration file is configuring the CapScan processor
# to be able to communicate with the CapScan server
#
# Capscan server name
server.host = 127.0.0.1
#
# Connection timeout in seconds (0 means no time out)
connection.timeout = 30
#
# The connection type to make to the CapScan server.
# Possible values are:
#
# CONNORIENTED
# CONNECTIONLESS
# STATELESS
```

```
# WEBCONNECTION
# ONDEMAND
connection.type = CONNECTIONLESS
#
# The number of connections the director server should
# make to the CapScan server
number.capscan.connections = 1
#
# The number of threads that should be used to
# communicate with the CapScan server
number.threads = 1
```

6

Using the Command Line Interface

This chapter describes how to use the command line interface. This chapter includes the following sections:

- [Running the Command Line Interface](#)
- [Understanding the Commands and Arguments](#)
- [Reviewing Examples](#)

The command line interface, `jmxttools.jar`, provides access to a number of facilities.

6.1 Running the Command Line Interface

The command line interface is distributed as a self contained `.jar` file in the `tools` directory, and is executed by the following command line invocation:

```
java -jar jmxttools.jar commandname arguments
```

The commands and arguments are described in the following section.

6.2 Understanding the Commands and Arguments

The command line interface can run a number of commands and provides functionality including:

- Running jobs
- Listing and dropping orphaned results tables
- Showing user session logs
- Shutting down real-time jobs
- Checking the EDQ version number

The following sections provide a full guide to the commands, arguments and options available.

EDQ also provides support for jobs through the REST-based EDQ Configuration API interfaces. For details on using these interfaces for performing various jobs, see "[REST Interfaces for Jobs](#)" in the "*Integrating Enterprise Data Quality With External Systems*" guide.

6.2.1 `runjob`

The `runjob` command runs a named job in the same way as if running the job using the Director UI. The `runjob` command takes the following arguments:

Argument	Use
<code>-job <i>job_name</i></code>	Specifies the name of the job to run.

Argument	Use
<code>-project project_name</code>	Specifies the name of the project that contains the job .
<code>-u user_name</code>	Specifies the user name to use to connect to the EDQ server. The user must have permission to run jobs and must have permission to the project containing the job.
<code>-p password</code>	Specifies the connecting user's password. If the <code>-p</code> option is not set, EDQ will prompt the user for the password.
<code>-nolockwait</code>	Indicates that if any of the resources used by the job are locked, the job should not wait for them to become available. Instead, it should terminate with a failure code and return control to the command line. The <code>-nolockwait</code> argument takes no extra values.
<code>-nowait</code>	Indicates that the command line should not wait for the job to complete. The <code>-nowait</code> argument takes no extra values.
<code>server:port</code>	Specifies the server and port of the JMX (management) interface.

6.2.2 `runopsjob`

The `runopsjob` command runs a named job in the same way as if running the job using the Server Console user interface. This provides additional functionality to the `runjob` command, specifically the use of Run Labels and Run Profiles. Run Labels may be used to store results separately from other runs of the same job. Run Profiles may be used to override externalized configuration settings at runtime.

The `runopsjob` command takes the following arguments:

Argument	Use
<code>-job job_name</code>	Specifies the name of the job to run.
<code>-project project_name</code>	Specifies the name of the project that contains the job
<code>-u user_name</code>	Specifies the user name to use to connect to the EDQ server. The user must have permission to run jobs and must have permission to the project containing the job.
<code>-p password</code>	Specifies the connecting user's password. If the <code>-p</code> option is not set, EDQ will prompt the user for the password.
<code>-nolockwait</code>	Indicates that if any of the resources used by the job are locked, the job should not wait for them to become available. Instead, it should terminate with a failure code and return control to the command line. The <code>-nolockwait</code> argument takes no extra values.
<code>-nowait</code>	Indicates that the command line should not wait for the job to complete. The <code>-nowait</code> argument takes no extra values.
<code>-runlabel run_label_name</code>	Specifies the name of the run label under which you wish to store staged output results. Note that this will override any run label that is specified in a run profile or by <code>-D runlabel = run_label_name</code> .
<code>-props run_profile_name</code>	Specifies the full path to a run profile properties file containing override settings for externalized configuration options in the job.

Argument	Use
<code>-D</code> <code>externalized_option</code> <code>=value</code>	Allows you to override specific externalized options for the job individually. The syntax for the externalized options and values is the same as used in run profile properties files. Note that characters is interpreted by the command line, so some characters will need to be escaped according to the shell conventions of your environment. Also note that any individually specified externalized option settings will override any settings for the same option if these are specified in a run profile used in the same run.
<code>server:port</code>	Specifies the server and port of the JMX (management) interface.

6.2.3 droporphans

The `droporphans` command is used to remove any orphaned results tables that may be created when processes are terminated unexpectedly. It should not be run when any jobs or processes are running on the EDQ server.

The `droporphans` command takes the following arguments:

Option	Use
<code>-u user name</code>	Specifies the user name to use to connect to the server. The user must have permission to cancel jobs and must have permission to the project containing the job.
<code>-p password</code>	Specifies the connecting user's password. If the <code>-p</code> option is not set, will prompt the user for the password.
<code>server:port</code>	Specifies the server and port of the JMX (management) interface.

6.2.4 listorphans

The `listorphans` command is used to identify any orphaned results tables. The `listorphans` command takes the same arguments as the `droporphans` command.

6.2.5 scriptorphans

The `scriptorphans` command creates a list of SQL commands for dropping orphaned results tables. This is useful if you want to review exactly which commands will run on the Results database when you drop tables, or if you want to drop the tables yourself manually.

6.2.6 list

The `list` command lists all the available commands.

6.2.7 showlogs

The `showlogs` command starts a small graphical user interface application that allows user session logs to be retrieved.

6.2.8 shutdown

The `shutdown` command shuts down all real-time jobs. These are jobs that are running from real-time record providers (web services or Java Message Service).

The `shutdown` command takes the following arguments:

Option	Use
<code>-u user name</code>	Specifies the user name to use to connect to the server. The user must have permission to cancel jobs and must have permission to the project containing the job.
<code>-p password</code>	Specifies the connecting user's password. If the <code>-p</code> option is not set, EDQ will prompt the user for the password.
<code>-nowait</code>	Indicates that the command line should not wait for the job to complete. The <code>-nowait</code> argument takes no extra values.
<code>server:port</code>	Specifies the server and port of the JMX (management) interface.

6.2.9 version

The `version` command is used to identify the version of the currently installed instance of .

Enter the following at the command line:

```
java -jar jmxtools.jar version
```

The version number is returned.

6.3 Reviewing Examples

This section lists several possible invocations of the command line interface:

- [Listing All the Available Commands](#)
- [Listing the Available Parameters for a Command](#)
- [Running a Named Job](#)
- [Running a Named Job in Operations Mode](#)

6.3.1 Listing All the Available Commands

The following invocation of the command line interface lists all of the available commands:

```
java -jar jmxtools.jar -list
```

The output is as follows:

```
Available launch names:
<Job tools>
runjob Run named job
shutdown Shutdown realtime jobs
runopsjob Run named job in operations mode
```

```
<Logging>
showlogs Show session logs

<Database Tools>
listorphans List orphaned results tables
droporphans Drop orphaned results tables
scriptorphans Create script for dropping orphaned results tables

<System Information>
version Display version number of tools
```

6.3.2 Listing the Available Parameters for a Command

If the command line interface is invoked by specifying a command without the corresponding parameters, it outputs detailed help for the command. For example, to get detailed help on the `runjob` command, invoke the command line interface as follows:

```
java -jar jmxtools.jar runjob
```

The output is as follows:

```
Usage: runjob -job jobname -project project [-u user] [-p pw] [-nowait] [-nolockwait]
[-sslprops props | -ssltrust store] server:port
```

6.3.3 Running a Named Job

This example illustrates how to run a named job in a named project on a specific instance (as specified by machine name and port).

To run a job called "rulecheck" in a project called "Audit" on the local machine with a JMX server on port 8090 using a user named "dnadmin", the command is as follows:

```
java -jar jmxtools.jar runjob -job rulecheck -project audit -u dnadmin
localhost:8090
```

The application prompts the user to enter the password for the `dnadmin` user.

6.3.4 Running a Named Job in Operations Mode

This example illustrates how to run a named job in 'operations mode' in a Windows environment. In operations mode, there is access to the Run Label and Run Profile capabilities so that the configuration of the job can be specified dynamically, and so that the results of the job can be stored by Run Label.

To run a job called "profiling" in a project called "MDM" on a server called "prod01", with a run label of "Nov2011" and a run profile file called `File1.properties`, with a JMX server on port 8090, the command is as follows:

```
java -jar jmxtools.jar runopsjob -job profiling -project MDM -runlabel Nov2011 -
props c:\ProgramData\Oracle\Enterprise Data
Quality\oedq_local_home\File1.properties" -u dnadmin prod01:8090
```

7

Configuring Additional Database Connections

This chapter describes how you can configure additional database connections for use in Director.

- [Using JNDI to Connect to Data Stores](#)
- [Connecting to an Oracle Database Using tnsnames.ora](#)
- [Connecting to an Oracle Database Using Oracle Internet Directory \(LDAP\)](#)

The standard options for Director to connect to data stores are described in the online help. Once implemented, these options appear in the Data Store Configuration step of the New Data Store wizard in Director. For help with using this wizard, see the Director online help.

7.1 Using JNDI to Connect to Data Stores

You can configure EDQ to use a Java Naming and Directory Interface (JNDI) data store connection.

1. Define the JNDI data store. JNDI is provided by the hosting application server. For more information about defining JNDI data sources in Oracle WebLogic Server, see "Using DataSource Resource Definitions" in .
2. In the EDQ data store wizard, specify JNDI as the type of data store, and then specify the JNDI name.

7.2 Connecting to an Oracle Database Using tnsnames.ora

You can configure EDQ to use an Oracle Transparent Network Substrate (TNS) data store connection. To use this connection method, you specify a name from a `tnsnames.ora` file as the data source when using the data sources wizard. Only the `tnsnames.ora` file is needed. No other Oracle client software is needed.

7.2.1 To Configure EDQ to Connect Through TNS

To connect EDQ through TNS:

1. Set the `oracle.net.tns_admin` Java system property to a local directory that contains the `tnsnames.ora` file.
2. Create a file named `jvm.properties` in your EDQ local configuration directory (`oedq_local_home` by default) and add an entry similar to the following:
`oracle.net.tns_admin = c:\\temp`. This property may have been set already in the application server when EDQ was installed.

For more information about the `tnsnames.ora` file, see "Configuring the Local Naming Method" in .

7.3 Connecting to an Oracle Database Using Oracle Internet Directory (LDAP)

You can configure EDQ to use an Oracle Lightweight Direct Access Protocol (LDAP) data store connection by setting the required Java system properties. These properties are:

```
dn.oracle.directory.servers = ldap://servername:port
```

```
dn.oracle.default.admin.context = dc=domaincontext1,dc=domaincontext2
```

The first property gives the location of your LDAP servers. The second property sets the context within the LDAP tree. Together, these properties enable EDQ to construct an Oracle and LDAP JDBC connection string, which looks similar to:

```
jdbc:oracle:thin:@ldap://servername:port/  
unicode,cn=Oraclecontext,dc=domaincontext1,dc=domaincontext2
```

8

Configuring EDQ to Process XML Data Files

This chapter describes how can be configured to read and write XML data files. This chapter includes the following sections:

- [Using Simple XML Data Stores](#)
- [Using XML and Stylesheet Data Stores](#)

You can use XML data files in snapshots to read and write the data contained in the file. A snapshot is a staged copy of data in a data store that is used in one or more processes. provides two types of data stores for working with XML data files: Simple XML and XML and Stylesheet. Both are available for server-side and client-side data stores.

8.1 Using Simple XML Data Stores

Simple XML data stores can read and write XML files that have a simple 2-level structure in which the top level tag represents the entity and the lower level tags represent the attributes of the entity. XML files exported from Microsoft Access are an example.

Following is an example of a simple XML file format that could be used with :

```
<dataroot>
  <Person>
    <Id>1</Id>
    <FirstName>Fred</FirstName>
    <LastName>Bloggs</LastName>
    <DateOfBirth>1972-01-31T00:00:00.000+0000</DateOfBirth>
    <Weight>85</Weight>
  </Person>
  <Person>
    <Id>2</Id>
    <FirstName>Jane</FirstName>
    <LastName>Smith</LastName>
    <DateOfBirth>1985-07-16T00:00:00.000+0100</DateOfBirth>
    <Weight>63</Weight>
  </Person>
</dataroot>
```

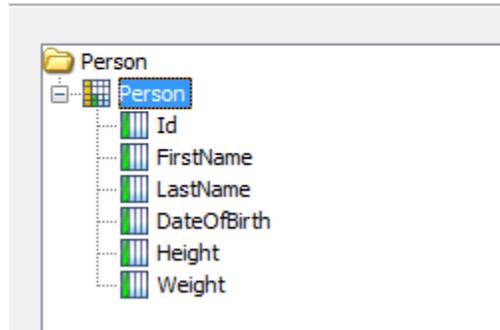
8.1.1 Reading Simple XML Files

When reads Simple XML files the following occurs:

- The root element name is not used, so it can be anything.
- The record element name appears as the table name in the Table Selection page of the Snapshot Wizard dialog.

Table Selection

What data do you want to snapshot?



- The lower level element names appear as the column names in the Column Selection page of the Snapshot Wizard and therefore become EDQ attribute names.

Column Selection

Which columns do you want to snapshot?

	Column Name	Data Type
<input checked="" type="checkbox"/>	Id	VARCHAR
<input checked="" type="checkbox"/>	FirstName	VARCHAR
<input checked="" type="checkbox"/>	LastName	VARCHAR
<input checked="" type="checkbox"/>	DateOfBirth	VARCHAR
<input checked="" type="checkbox"/>	Height	VARCHAR
<input checked="" type="checkbox"/>	Weight	VARCHAR

8.1.2 Writing Simple XML Files

When generating Simple XML files using an export to the data store, the name of the data store defines the record XML element name. The element `Person` in the example in [Using Simple XML Data Stores](#) shows how this appears in the XML.

The XML element names of the lower level tags are taken from the attribute names. EDQ names are encoded to ensure that invalid XML is not generated. For example, space characters in names are replaced by the character sequence `_x0020_`, so an attribute named `Date Of Birth` would generate XML elements in the following format:

```
<Date_x0020_Of_x0020_Birth>
```

8.2 Using XML and Stylesheet Data Stores

When there is a requirement to work with XML of a different structure than that of Simple XML, then you use the XML and Stylesheet data stores.

These data stores read and write XML conforming to the DN-XML schema and optionally allow the use of a custom stylesheet to:

- Transform XML from a custom XML format to DN-XML during data snapshot
- Transform XML from DN-XML to a custom XML format during data export

For more information about XML stylesheets, see the W3C website found at <http://www.w3.org/Style/XSL/> and <http://www.w3.org/standards/xml>.

8.2.1 Using DN-XML

DN-XML is the format by which custom XML can be processed by .

An example of DN-XML is as follows:

```
<dn:data xmlns:dn="http://www.datanomic.com/2008/dnx">
  <dn:record skip="true">
    <dn:value name="Id" type="string"/>
    <dn:value name="FirstName" type="string"/>
    <dn:value name="LastName" type="string"/>
    <dn:value name="DateOfBirth" type="date"/>
    <dn:value name="Height" type="number"/>
    <dn:value name="Weight" type="number"/>
  </dn:record>
  <dn:record>
    <dn:value name="Id">1</dn:value>
    <dn:value name="FirstName">Fred</dn:value>
    <dn:value name="LastName">Bloggs</dn:value>
    <dn:value name="DateOfBirth">1972-01-31</dn:value>
    <dn:value name="Height">1.85</dn:value>
    <dn:value name="Weight">85</dn:value>
  </dn:record>
  <dn:record>
    <dn:value name="Id">2</dn:value>
    <dn:value name="FirstName">Jane</dn:value>
    <dn:value name="LastName">Smith</dn:value>
    <dn:value name="DateOfBirth">1985-07-16</dn:value>
    <dn:value name="Height">1.65</dn:value>
    <dn:value name="Weight">63</dn:value>
  </dn:record>
</dn:data>
```

This is the equivalent DN-XML for the example given in [Using Simple XML Data Stores](#).

Note that the attribute names are defined differently in DN-XML compared with Simple XML. Because DN-XML uses attribute content to specify attribute names, it is possible to create attributes with spaces and other special characters in their names.

In the previous example, the `<dn:record skip="true">` XML element and its contents allows the definition of the structure of the source including the field names and their data types. All other record elements define a row of data in . This is analogous to the header row in a comma-separated values file. The following data types are permitted:

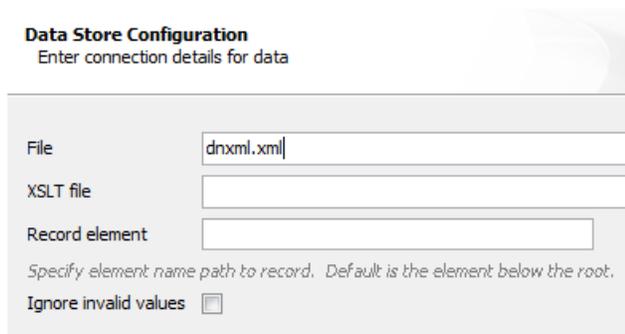
- string
- date
- number

 **Note:**

Date values in DN-XML files should be specified in the XSD date format (ISO 8601). For example, '2008-10-31T15:07:38.6875000-05:00' or without the time component simply as '2008-10-31'.

Within a data record, value elements are used to specify attribute values for the record. The name attribute is used to specify the attribute in question and the text content of the attribute specifies the value for that attribute. For example, the XML fragment, `<dn:value name="FirstName">Fred</dn:value>`, assigns the value 'Fred' to the attribute 'FirstName'.

DN-XML files can be read in to by creating an XML and Stylesheet data store and specifying the location of the XML source file; the XSLT file option should be left blank:



Data Store Configuration
Enter connection details for data

File:

XSLT file:

Record element:

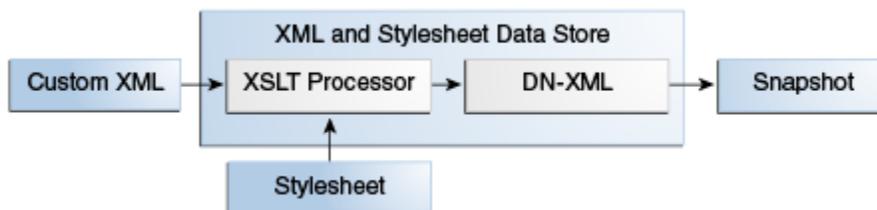
Specify element name path to record. Default is the element below the root.

Ignore invalid values

Similarly, can write DN-XML files by exporting data to an XML and Stylesheet data store with the XSLT option left blank.

8.2.2 Reading Custom XML Files

XML files in custom formats can be read by using the XML and Stylesheet data store configured to use a custom XML stylesheet (XSLT) to transform from the custom schema to the DN-XML schema during data snapshotting.



Following is an example custom XML file that could be read into :

```

<crmdata>
  <contacts>
    <contact id="1">
      <name>
    
```

```

        <firstname>Fred</firstname>
        <surname>Bloggs</surname>
    </name>
    <dob>1972-01-31</dob>
    <properties>
        <property name="height" value="1.85"/>
        <property name="weight" value="85"/>
    </properties>
</contact>
<contact id="2">
    <name>
        <firstname>Jane</firstname>
        <surname>Smith</surname>
    </name>
    <dob>1985-07-16</dob>
    <properties>
        <property name="height" value="1.68"/>
        <property name="weight" value="63"/>
    </properties>
</contact>
</contacts>
</crmdata>

```

The following XML stylesheet demonstrates one way that the preceding example custom XML can be transformed into a suitable DN-XML format:

```

<xsl:stylesheet version="1.0" xmlns:dn="http://www.datanomic.com/2008/dnx"
  xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
  xmlns:fn="http://www.w3.org/2005/02/xpath-functions">

  <xsl:output method="xml"/>

  <xsl:template match="/">
    <dn:data>

      <!-- Write out the header record -->
      <dn:record skip="true">
        <dn:value name="Id" type="string"/>
        <dn:value name="FirstName" type="string"/>
        <dn:value name="LastName" type="string"/>
        <dn:value name="DateOfBirth" type="date"/>
        <dn:value name="Height" type="number"/>
        <dn:value name="Weight" type="number"/>
      </dn:record>

      <!-- Get each contact record -->
      <xsl:apply-templates select="/crmdata/contacts/contact"/>

    </dn:data>
  </xsl:template>

  <xsl:template match="contact">

    <!-- Write out a data record -->
    <dn:record>
      <dn:value name="Id"><xsl:value-of select="@id"/></dn:value>
      <dn:value name="FirstName"><xsl:value-of select="name/firstname"/></dn:value>
      <dn:value name="LastName"><xsl:value-of select="name/surname"/></dn:value>
      <dn:value name="DateOfBirth"><xsl:value-of select="dob"/></dn:value>
      <dn:value name="Height">

```

```

    <xsl:value-of select="properties/property[@name='height']/@value"/>
  </dn:value>
  <dn:value name="Weight">
    <xsl:value-of select="properties/property[@name='weight']/@value"/>
  </dn:value>
</dn:record>

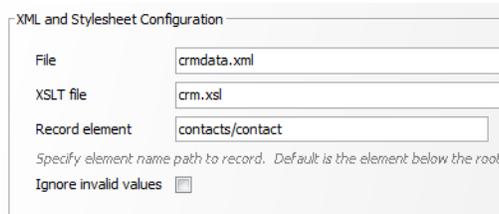
</xsl:template>

</xsl:stylesheet>

```

8.2.2.1 Configuring the Data Store

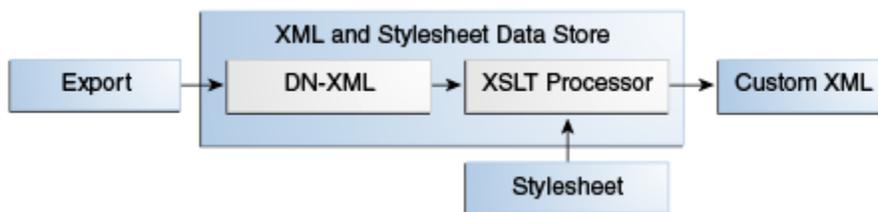
The data can be read in to by creating an XML and Stylesheet data store and specifying the location of the XML source file and the XSLT file (stylesheet).



reads the source XML file in chunks for efficiency breaking up the file on record boundaries. By default uses the element immediately below the root as the record element. If this is not the case in the source XML file then an XPath-style expression to the record element from the root must be specified.

8.2.3 Writing Custom XML Files

XML files in custom formats can be written by using the XML and Stylesheet data store configured to use a custom XSLT to transform from the DN-XML schema to the custom target schema the during data export.



Following is an example target custom XML format that needs to be generated by :

```

<Report>
  <Person Id="1" FullName="Fred Bloggs"/>
  <Person Id="2" FullName="Jane Smith"/>
</Report>

```

The following XML stylesheet demonstrates one way in which the DN-XML format can be transformed into the target custom XML format:

```
<xsl:stylesheet version="1.0"
  xmlns:dn="http://www.datanomic.com/2008/dnx"
  xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
  xmlns:fn="http://www.w3.org/2005/02/xpath-functions">

  <xsl:output method="xml"/>

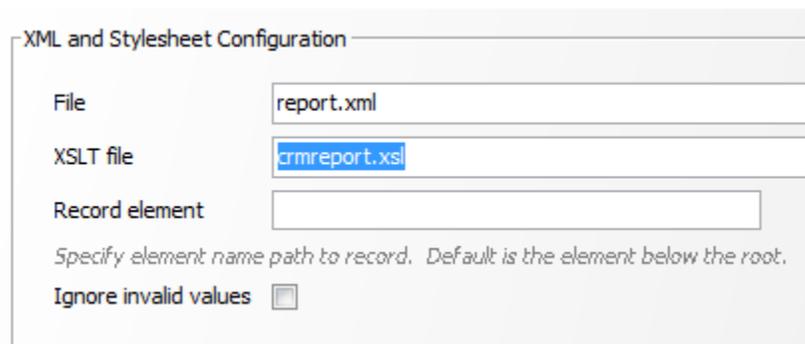
  <xsl:template match="/">
    <Report>
      <xsl:apply-templates select="/dn:data/dn:record"/>
    </Report>
  </xsl:template>

  <xsl:template match="dn:record">
    <Person>
      <xsl:attribute name="Id">
        <xsl:value-of select="dn:value[@name = 'Id']"/>
      </xsl:attribute>
      <xsl:attribute name="FullName">
        <xsl:value-of select="dn:value[@name = 'FirstName']"/>
        <xsl:text> </xsl:text>
        <xsl:value-of select="dn:value[@name = 'LastName']"/>
      </xsl:attribute>
    </Person>
  </xsl:template>

</xsl:stylesheet>
```

8.2.3.1 Configuring the Data Store

The data can be written by by creating an XML and Stylesheet data store and specifying the destination for the custom XML file and XSLT (stylesheet) file.



The screenshot shows a configuration window titled "XML and Stylesheet Configuration". It contains the following fields and options:

- File:** report.xml
- XSLT file:** crmreport.xsl
- Record element:** (empty text box)
- Specify element name path to record. Default is the element below the root.*
- Ignore invalid values:**

9

Using the EDQ Configuration API

EDQ provides a set of REST-based interfaces that enable you to perform various configuration tasks programmatically, using any preferred programming language. In this chapter, the EDQ service is assumed to be installed at:

```
http://edqserver:8001/edq
```

This chapter provides a detailed description of these interfaces and the operations that can be performed using these interfaces. It includes the following topics:

- [REST Interface for Projects](#)
- [REST Interface for Data Stores](#)
- [REST Interface for Snapshots](#)
- [REST Interface for Processes](#)
- [REST Interfaces for Jobs](#)
- [REST Interface for Reference Data](#)
- [REST Interface for Web Services](#)
- [Example: Profiling from an External Application](#)

9.1 REST Interface for Projects

The REST interface for working with EDQ projects is

```
http://edqserver:8001/edq/config/projects
```

This interface allows you to perform the following tasks:

- [Retrieving a List of EDQ Projects](#)
- [Creating a Project](#)
- [Deleting a Project](#)

9.1.1 Retrieving a List of EDQ Projects

To get a list of all projects that are available with the current EDQ installation, you need to simply run an HTTP GET operation on the REST interface for EDQ projects, as shown in the following code:

```
GET http://edqserver:8001/edq/config/projects
```

When this code runs successfully, a list of projects is generated in JSON format:

```
[
  {
    "id":10,
    "name":"My Project"
  },

```

```
{
  "id":12,
  "name":"Scratch Project"
}
```

9.1.2 Creating a Project

To create a new project, you need to create a JSON object that describes the project to be created and then send it in the request body of the REST call using an HTTP POST.

For example:

```
POST http://edqserver:8001/edq/config/projects
```

```
{ "name" : "Profile Customer Names" ,
  "description" : "Profile my customers" }
```

This code returns a response of type "OK" with a response body similar to the following:

```
{"id":14,"name":"Profile Customer Names"}
```

However, if an error occurs while creating a project, then a response of type "500 Internal Server Error" is generated, along with an error message similar to the following:

```
"Profile Customer Names" already exists (Code: 205,130)
```

9.1.3 Deleting a Project

To delete a project, you need to call HTTP DELETE on the REST interface and specify the project you need to delete, as a query parameter.

There are two ways to specify a project for deletion:

- By ID using `pid=<NN>`
- By name using `pname=<Name>`

To delete a project by ID, do the following:

```
DELETE http://edqserver:8001/edq/config/projects?pid=14
```

To delete a project by name, do the following:

```
DELETE http://edqserver:8001/edq/config/projects?pname=Profile%20Customer%20Names
```

In both cases, the result is a string similar to the following:

```
Project Profile Customer Names deleted
```

If an invalid project is specified then a response of type 406 'Not Acceptable' is returned with an appropriate string message, for example:

```
Bad project ID "14" (Code: 205,454)
```

or

```
No project named "Profile Customer Names" (Code: 205,453)
```

9.2 REST Interface for Data Stores

You can query and manipulate data stores in EDQ, using the following interface:

```
http://edqserver:8001/edq/config/datasources
```

This interface allows you to perform the following tasks:

- [Retrieving a List of Data Stores](#)
- [Creating a Data Store](#)
- [Deleting a data store](#)

9.2.1 Retrieving a List of Data Stores

To get a list of data stores, call the interface with a valid project name.

For example:

```
GET http://edqserver:8001/edq/config/datasources?pid=14
```

If successful, an OK response is returned along with a list of data stores in the response body.

For example:

```
[
  {
    "client":false,
    "id":36,
    "name":"Individuals",
    "properties":[
      {
        "name":"quote",
        "value":"\"\"\"
      },
      {
        "name":"encoding",
        "value":"ISO-8859-1"
      },
      {
        "name":"file",
        "value":"Customer/customerindividuals.csv"
      },
      {
        "name":"cols",
        "value":""
      },
      {
        "name":"project",
        "value":"59"
      },
      {
        "name":"hdr",
        "value":"1"
      },
      {
        "name":"usepr",
        "value":"0"
      }
    ]
  }
]
```

```

    },
    {
      "name": "skip",
      "value": ""
    },
    {
      "name": "sep",
      "value": ", "
    }
  ],
  "species": "servertxt"
}
]

```

9.2.2 Creating a Data Store

To create a data store you need to create a JSON object that describes the data store and then POST it to the endpoint specifying the project (by name or id) that will own the data store.

For example:

A server-based .csv file that has been placed in the landing area.

POST <http://edqserver:8001/edq/config/datasources?pid=14>

```

{
  "client": false,
  "name": "Individuals",
  "properties": [
    {
      "name": "quote",
      "value": "\""
    },
    {
      "name": "encoding",
      "value": "ISO-8859-1"
    },
    {
      "name": "file",
      "value": "Customer/customerindividuals.csv"
    },
    {
      "name": "hdr",
      "value": "1"
    },
    {
      "name": "usepr",
      "value": "0"
    },
    {
      "name": "sep",
      "value": ", "
    }
  ],
  "species": "servertxt"
}

```

A server-based Oracle schema:

POST http://edqserver:8001/edq/config/datasources?pid=14

```
{
  "client":false,
  "name":"Staging",
  "properties":[
    {
      "name":"service",
      "value":"sid"
    },
    {
      "name":"sid",
      "value":"orcl"
    },
    {
      "name":"user",
      "value":"staging"
    },
    {
      "name":"port",
      "value":"1521"
    },
    {
      "name":"password",
      "value":"staging"
    },
    {
      "name":"host",
      "value":"localhost"
    }
  ],
  "species":"oracle"
}
```

If successful, an OK response is returned along with the name and ID of the data store, as shown in the following example:

```
{"id":42,"name":"Staging"}
```

The value of the `species` parameter varies depending on the type of data store being used in a project. For example, if you are using the Oracle database, the value of `species` would be "oracle". Each species parameter has its own set of properties.

The following table lists the properties for the species "oracle":

Property	Type	Required	Description
host	String	Yes	The machine hosting the database
port	Number	Yes	Port number of the database
sid	String	Yes	Database Identifier
service	Choice of sid or srv	Yes	Whether the database identifier above is a SID or a SERVICE (srv) name
user	String	Yes	User to log in as
password	String	No	Password for the user
schema	String	No	The schema to use (usually left empty)

The following table lists the properties for the species "servertext":

Property	Type	Required	Description
file	String	Yes	The name and location of the file in the landing area
userpr	Boolean	Yes	Use project specific landing area
hdr	Boolean	Yes	Treat the first line as header
sep	String	No	The field delimiter
quote	Choice	No	Quote character
cols	Integer	No	Number of columns to read
encoding	String	Yes	The character encoding of text
skip	Integer	No	The number of the lines to skip at the start

 **Note:**

For Boolean type use 0 for false and 1 for true.

For quote, the value needs to be a double quote like this ""\", or a single quote like this '\', or an empty value like this ''.

For the species, "other", which uses a JDBC connection, the properties are listed in the following table:

Property	Type	Required	Description
driver	String	Yes	The JDBC driver java class
url	String	Yes	The address of the database
user	String	No	The user to log in as
password	String	No	The user's password

9.2.3 Deleting a data store

To delete a data store call HTTP DELETE on the endpoint by specifying either a data store id or a valid project (by name or id) and a data store name, for example:

```
DELETE http://edqserver:8001/edq/config/datasources?id=42
```

or

```
DELETE http://edqserver:8001/edq/config/datasources?pid=14&name=Staging
```

When the deletion is successful, an OK response is returned without any response body.

9.3 REST Interface for Snapshots

The REST interface for snapshots is:

`http://edqserver:8001/edq/config/snapshots`

It allows you to perform the following tasks:

- [Retrieving a List of Snapshots](#)
- [Creating a Snapshot](#)
- [Deleting a Snapshot](#)

9.3.1 Retrieving a List of Snapshots

To retrieve a list of snapshots specify a valid project, for example:

```
GET http://edqserver:8001/edq/config/snapshots?pid=14
```

or

```
GET http://edqserver:8001/edq/config/snapshots?pname=Profile%20Customer%20Names
```

If successful, it returns an OK response with a list of snapshots in the response body.

For example:

```
[
  {
    "columns": [
      "TITLE",
      "FULLNAME",
      "GIVENNAMES",
      "FAMILYNAME",
      "NAMETYPE",
      "PRIMARYNAME",
      "ADDRESS1",
      "ADDRESS2",
      "ADDRESS3",
      "ADDRESS4",
      "CITY",
      "STATE",
      "POSTALCODE"
    ],
    "datasource": "Individuals",
    "name": "Individuals",
    "table": "customerindividuals.csv"
  }
]
```

9.3.2 Creating a Snapshot

To create a snapshot, you need to create a JSON object that describes the snapshot and specify the project where it will be created.

For example:

```
POST http://edqserver:8001/edq/config/snapshots?pid=14

{
  "name": "Individuals",
  "description": "Customer data",
  "datasource": "Individuals",
  "table": "customerindividuals.csv",
```

```

    "columns":[
      "TITLE",
      "FULLNAME",
      "GIVENNAMES",
      "FAMILYNAME",
      "NAMETYPE",
      "PRIMARYNAME",
      "ADDRESS1",
      "ADDRESS2",
      "ADDRESS3",
      "ADDRESS4",
      "CITY",
      "STATE",
      "POSTALCODE"
    ],
    "sampling":{
      "number":100,
      "offset":0,
      "ordering":"ascending",
      "count":"true"
    }
  }
}

```

If successful, an OK response is returned with the snapshot ID in the response body.

For example:

```
{"id":68,"name":"Individuals"}
```

9.3.3 Deleting a Snapshot

To delete a snapshot, you need to specify either a snapshot ID or a valid project and snapshot name.

For example:

```
DELETE http://edqserver:8001/edq/config/snapshots?id=68
```

or

```
DELETE http://edqserver:8001/edq/config/snapshots?pid=14&name=Individuals
```

or

```
DELETE http://edqserver:8001/edq/config/snapshots?
pname=Profile%20Customer%20Names&name=Individuals
```

When the specified snapshot is deleted successfully, it returns an OK response with a string message in the response body, which is similar to the following:

```
Snapshot Individuals deleted
```

9.4 REST Interface for Processes

The interface for EDQ processes is:

```
http://edqserver:8001/edq/config/processes
```

Using this interface, you can perform the following tasks:

- [Retrieving a List of Processes](#)
- [Deleting a Process](#)

The following sub-level interface allows you to create a simple profiling process:

```
http://edqserver:8001/edq/config/processes/simpleprocess
```

See [Creating a Simple Process](#) section for details.

9.4.1 Retrieving a List of Processes

To get a list of processes in a project, you need to call HTTP GET on the processes interface with a valid project name.

Example:

```
GET http://edqserver:8001/edq/config/processes?pid=14
```

or

```
GET http://edqserver:8001/edq/config/processes?pname=Profile%20Customer%20Names
```

An OK response is returned along with a list of processes.

Example:

```
[{"name":"Profile Names","id":31}]
```

If the request is not successful, an error response would either be a 404 'Not Found' or 500 'Internal Server Error' along with a string in the response body describing the error.

9.4.2 Deleting a Process

To delete a process, specify either the process ID, or the project ID or name, and the process name. For example:

```
DELETE http://edqserver:8001/edq/config/processes?id=31
```

or

```
DELETE http://edqserver:8001/edq/config/processes?pid=14&name=Profile%20Names
```

or

```
DELETE http://edqserver:8001/edq/config/processes?  
pname=Profile%20Customer%20Names&name=Profile%20Names
```

When the deletion is successful, an OK response is returned with a string message and response body, as shown in the following example:

```
Process Profile Names deleted
```

If the deletion is not successful, then either of the following errors along with a response string are returned:

- 404 "Not Found"
- 500 "Internal Server Error"

9.4.3 Creating a Simple Process

The interface currently only supports creation of simple profiling processes. To create a simple process, you need to create a JSON object that describes the process you want to create and specify the project where it should be created. The HTTP POST operation is used to post this information to the interface.

For example:

```
POST http://edqserver:8001/edq/config/processes/simpleprocess?pid=14
{
  "name":"Profile Names",
  "description":"Profile Individuals Names",
  "reader":{
    "name":"Read from Individuals",
    "stageddata":"Individuals"
  },
  "processors":[
    {
      "name":"Do Quickstats",
      "type":"dn:quickstatsprofiler",
      "columnlist":[
        "GivenNames",
        "FamilyName"
      ]
    },
    {
      "name":"Do Frequency Profiling",
      "type":"dn:attributefrequencycountsprofiler"
    }
  ]
}
```

When the simple process is created successfully, an OK response is returned along with the name and ID of the process in the response body.

For example:

```
{"id":33,"name":"Profile Names"}
```

An error response would be generated in the following cases:

- 404 'Not Found' if the project does not exist
- 400 'Bad Request' if the JSON object is malformed
- 500 'Internal Server Error' if a server error occurs during creation, along with a string in the response body describing the error.

The full list of supported processors is mentioned in the following table:

Processor	Type
Quick Stats Profiler	dn:quickstatsprofiler
Data Types Profiler	dn:datatypesprofiler
Max/Min Profiler	dn:maxandminprofiler
Length Profiler	dn:lengthprofiler

Processor	Type
Record Completeness Profiler	dn:recordcompletenessprofiler
Character Profiler	dn:characterprofiler
Frequency Profiler	dn:attributefrequencycountsprofiler
Patterns Profiler	dn:attributepatternsprofiler

9.5 REST Interfaces for Jobs

The `/config/jobs` interface performs the following tasks on EDQ jobs:

- [Retrieving a List of Jobs](#)
- [Deleting a Job](#)
- [Creating a Simple Job](#)

The URL for this interface is similar to the following:

```
http://edqserver:8001/edq/config/jobs
```

There is another REST interface, `/jobs` to perform the following tasks:

- [Running a Job](#)
- [Cancelling a Running Job](#)
- [Getting the Status of a Job](#)
- [Getting the Details of All Running Jobs](#)

The URL for this interface is similar to the following:

```
http://edqserver:8001/edq/jobs
```

9.5.1 Retrieving a List of Jobs

You can get a list of jobs for a project using HTTP GET. You must specify at least one project (there could be more than one project) as the query parameter.

The project can be specified by using the `pid` or `pname` in the query parameter, as shown in the following example:

```
GET http://edqserver:8001/edq/config/jobs?pid=14
```

or

```
GET http://edqserver:8001/edq/config/jobs?pname=Profile%20Customer%20Names
```

The response to this operation would list all the jobs for the specific project or projects, as shown in the following example:

```
[
  {
    "id":99,
    "name":"Profile Names Job"
  },
  {
    "id":98,
```

```
        "name": "Profile Individuals Job"
    }
]
```

9.5.2 Deleting a Job

To delete a job you need to either specify a valid job ID or a valid project (using one of the query parameters) and a valid job name.

To delete a job by job ID:

```
DELETE http://edqserver:8001/edq/config/jobs?id=99
```

To delete a job by job name:

```
DELETE http://edqserver:8001/edq/config/jobs?
pid=14&name=Profile%20Names%20Job
```

or

```
DELETE http://edqserver:8001/edq/config/jobs?
pname=Profile%20Customer%20Names&name=Profile%20Names%20Job
```

When the job is deleted successfully, a message appears to confirm the deletion:

```
Job Profile Names Job deleted
```

9.5.3 Creating a Simple Job

A simple job has a single phase and contains a single process. To create a simple job, you need to specify a valid project that owns the job (using one of the query parameters). Also, you need to create a JSON object describing the job to create.

For example:

```
POST http://edqserver:8001/edq/config/jobs/simplejob?pid=14
{
  "name" : "Profile Names Job" ,
  "process" : "Profile Names" ,
  "description" : "Profile Customer Names"
  "resultsdrilldown" : "none"
}
```

The attribute `resultsdrilldown` can have any of the following values: `none`, `sample`, `limited`, `all`. However, the `sample` and `limited` values have the same implication.

9.5.4 Running a Job

With the `/jobs/run` interface you can run a named job using HTTP POST. The required parameters are the project name or project ID and the job name. Optionally, you can specify run label and overrides.

The following example shows the URL and associated payload used with running a job named "Real-time Start All" for the project "Profile Customer Name":

```
POST http://edqserver:8001/edq/jobs/run
{
  "project": "Profile Customer Name",
```

```
"job": "Real-time Start All",
"runlabel": "uk",
"runprofile": "profiling",
"overrides": [
  {
    "name": "a",
    "value": "b"
  },
  {
    "name": "c",
    "value": "d"
  }
]
```

The JSON response to this request would be similar to the following:

```
{
  "executionID": 2,
  "runeverywhere": false
}
```

In this example, the job "Real-time Start All" returns the "runeverywhere" value as false, which implies that this job can run only in one place. In such cases, an `executionID` is returned for the job. This `executionID` can be used to cancel a job and query a job's status.

However, if the value of "runeverywhere" were true, then only the "jobtype" would be returned in the JSON response. For "runeverywhere" jobs, cancel and query calls are not supported.

9.5.5 Cancelling a Running Job

To cancel a running job, the HTTP POST operation is used. The interface URL is similar to the following:

```
POST http://edqserver:8001/edq/jobs/cancel
```

You only need the `executionID` of the job to cancel it. The following example illustrates cancelling a job with the `executionID` 12.

```
{
  "executionID": 12,
  "type" : "immediate"
}
```

The following options are available with the "type" parameter:

- `immediate`: This option cancels the job as early as possible.
- `keepresults`: This option cancels the job but retains the results that have been generated so far.
- `shutdown`: This options is used to cancel or shutdown a job that runs a web service.

No response is returned when a job is cancelled successfully.

9.5.6 Getting the Status of a Job

To get the status of an individual job, the `executionID` is passed in the URL. The URL looks similar to the following:

```
GET http://edqserver:8001/edq/jobs/status?xid=executionID
```

For example, if the execution ID of the job "Real-time Start All" is 14, the URL would be:

```
GET http://edqserver:8001/edq/jobs/status?xid=14
```

The JSON response would be as follows:

```
{
  "executionid": 14,
  "project": "Profile Customer Name",
  "job": "Real-time START ALL",
  "server": "edqserver",
  "starttime": "2016-04-20T08:44:48+01:00",
  "endtime": "2016-04-20T08:45:22+01:00",
  "complete": true,
  "status": "finished"}
```

In this example, the "Real-time Start All" job triggers other jobs. Once all the jobs in the project are triggered, the status of the `executionID` 14 shows finished. However, the jobs that are triggered by the "Real-time Start All" job, may still show the status as running.

9.5.7 Getting the Details of All Running Jobs

The status of all running jobs in a project can be retrieved using the `/jobs/running` interface, which would be represented by a URL similar to the following:

```
GET http://edqserver:8001/edq/jobs/running
```

The following example shows the output in JSON format:

```
{
  "executionid": 4,
  "project": "Profile Customer Name",
  "job": "Real-time Individual Clean",
  "server": "edqserver",
  "starttime": "2016-04-19T10:05:30.74+01:00",
  "endtime": "2016-04-19T16:05:41+01:00",
  "complete": false,
  "status": "running"
}
```

Optionally, you can provide other query parameters such as project name, job name, and run label. For jobs without a run label, omit the run label parameter and set it to empty filter jobs with no run label. The URL in this case would be similar to the following:

```
/jobs/running?[project=project[&job=job][&runlabel=]
```

9.6 REST Interface for Reference Data

Reference data can exist within a project or outside of all projects at system level. To refer to reference data at system level, specify the project ID as 0 or `pid=0`.

The interface for reference data is:

```
http://edqserver:8001/edq/config/referencedata
```

The interface for reference data content is:

```
http://edqserver:8001/edq/config/referencedata/contents
```

You can use this interface to perform the following tasks:

- [Retrieving a List of Reference Data](#)
- [Retrieving Contents of Reference Data](#)
- [Creating Reference Data](#)
- [Deleting Reference Data](#)

9.6.1 Retrieving a List of Reference Data

To get a list of all reference data defined at system level specify the project using either of the following parameters:

By `pid = 0`:

```
GET http://edqserver:8001/edq/config/referencedata?pid=0
```

By `pname`:

```
GET http://edqserver:8001/edq/config/referencedata?  
pname=Profile%20Customer%20Names
```

A list of JSON objects, which represent reference data, are returned. The output looks similar to the following:

```
[  
  {  
    "activerows":2,  
    "category":"characterkeymap",  
    "columns": [  
      {  
        "key":true,  
        "name":"Name",  
        "type":"STRING",  
        "unique":true  
      },  
      {  
        "name":"Value",  
        "type":"STRING",  
        "value":true  
      }  
    ],  
    "id":39,  
    "name":"Tokens",  
    "totalrows":2  
  }  
]
```

```
    }  
  ]  
}
```

9.6.2 Retrieving Contents of Reference Data

To list the contents of the reference data, the reference data contents interface is used. You must specify a valid project or use `pid=0` for system level.

For example:

```
GET http://edqserver:8001/edq/config/referencedata/contents?pid=0&id=40
```

or

```
GET http://edqserver:8001/edq/config/referencedata/contents?  
pid=0&name=ShortNameMap
```

This returns information about the reference data rows, as shown in the following code snippet:

```
{  
  "activerows":4,  
  "columns":[  
    {  
      "key":true,  
      "name":"ShortName",  
      "type":"STRING",  
      "unique":true,  
      "value":true  
    },  
    {  
      "key":true,  
      "name":"LongName",  
      "type":"STRING",  
      "value":true  
    }  
  ],  
  "description":"Map short names to long names",  
  "id":43,  
  "name":"ShortNameMap",  
  "rows":[  
    {  
      "data":[  
        "Jeff",  
        "Jeffrey"  
      ]  
    },  
    {  
      "data":[  
        "Jon",  
        "Jonathan"  
      ]  
    }  
  ],  
  "totalrows":4  
}
```

9.6.3 Creating Reference Data

To create reference data you need to create a JSON object describing the reference data, which you will post to the interface specifying either `pid=0` for system level, or a valid project name.

For example:

```
POST
http://edqserver:8001/edq/config/referencedata?pname=Profile%20Customer%20Names
```

```
{
  "name" : "ShortNameMap",
  "description" : "Map short names to long names",
  "columns":
    [
      {
        "key": true,
        "name": "ShortName",
        "type": "STRING",
        "unique": true,
        "value": true
      },
      {
        "key": true,
        "name": "LongName",
        "type": "STRING",
        "value": true
      }
    ],
  "rows":
    [
      {
        "data":
          [
            "Jeff",
            "Jeffrey"
          ]
      },
      {
        "data":
          [
            "Jon",
            "Jonathan"
          ]
      }
    ]
}
```

On successful creation, a response similar to the following is returned:

```
{"id":40,"name":"ShortNameMap"}
```

9.6.4 Deleting Reference Data

To delete reference data, you need to call HTTP DELETE on the reference data interface, with either a valid reference data ID or a valid project (including `pid=0` for system level) and a valid reference data name.

To delete by reference data ID:

```
DELETE http://edqserver:8001/edq/config/referencedata?id=40
```

To delete by reference data name:

```
DELETE http://edqserver:8001/edq/config/referencedata?pid=0&name=ShortNameMap
```

After a successful deletion, the response returns a string message, such as the following:

```
Reference data ShortNameMap deleted
```

9.7 REST Interface for Web Services

The interface for web services is:

```
http://edqserver:8001/edq/config/webservices
```

It allows you to perform the following tasks when you call the respective get, post, or delete operations:

- [Retrieving a List of Web Services](#)
- [Creating or Updating a Web Service](#)
- [Deleting a Web Service](#)

9.7.1 Retrieving a List of Web Services

You can get a list of web services defined for a valid project by using the HTTP GET operation on the web services interface. To get a list of web services, specify a valid project using either the `pid` or `pname` parameter.

For example:

```
GET http://edqserver:8001/edq/config/webservices?pid=14&pid=20
```

A successful call returns a list of web services, and their input and output interfaces, in the response body:

```
[
  {
    "id":1,
    "inputs":{
      "attributes":[
        {
          "name":"Name",
          "type":"STRING"
        }
      ],
      "multirecord":false
    },
    "name":"Long Names",
    "outputs":{
      "attributes":[
        {
          "name":"LongName",
          "type":"STRING"
        }
      ],

```

```

        "multirecord":false
    }
}
]

```

9.7.2 Creating or Updating a Web Service

You can create and update web services by creating an appropriate JSON object, which you then POST to the web services interface.

To create a web service you need to specify a valid project (by name or id), as shown in the following example:

POST <http://edqserver:8001/edq/config/webservices?pid=14>

```

{
  "name":"Name Gender",
  "inputs":{
    "attributes":[
      {
        "name":"Name",
        "type":"STRING"
      }
    ],
    "multirecord":false
  },
  "outputs":{
    "attributes":[
      {
        "name":"Gender",
        "type":"STRING"
      }
    ],
    "multirecord":false
  }
}

```

If successful, the name and ID of the web service is returned in the response body.

Example:

```
{"id":4,"name":"Name Gender"}
```

To update a web service, you need a JSON object that is identical in structure, but with an additional ID attribute to identify the existing web service. For an update you do not specify a project.

Example:

POST <http://edqserver:8001/edq/config/webservices>

```

{
  "id":4,
  "name":"Name Gender",
  "inputs":{
    "attributes":[
      {
        "name":"First Name",
        "type":"STRING"
      },
      {

```

```

        "name": "Last Name",
        "type": "STRING"
      }
    ],
    "multirecord": false
  },
  "outputs": {
    "attributes": [
      {
        "name": "Gender",
        "type": "STRING"
      }
    ],
    "multirecord": false
  }
}

```

If successful, the name and ID of the web service is returned in the response body, as shown in the following example:

```
{"id":4,"name":"Name Gender"}
```

9.7.3 Deleting a Web Service

To delete a web service call HTTP DELETE on the web service interface. Specify either the web service ID or a valid project (by name or ID) and the web service name.

For example:

```
DELETE http://edqserver:8001/edq/config/webservices?id=5
```

or

```
DELETE http://edqserver:8001/edq/config/webservices?
pid=14&name=Name%20Gender
```

If successful, an OK response is returned but without a response body.

9.8 Example: Profiling from an External Application

Consider a scenario where an external application needs to profile data in a table in an Oracle database, using EDQ. In such a case, you can programmatically profile this table using the REST-based APIs. For this example, a CUSTOMERS table in a CustomerDB database will be used.

To generate and run the profiling job on the CUSTOMERS table, the following tasks are performed:

1. Create a project by using the following URL:

```
POST http://edqserver:8001/edq/config/projects
```

The project name (`pname`) is "Profile Customer". The JSON code is:

```

{
  "name": "Profile Customer",
  "description": "Profiling customers in the CUSTOMERS table"
}

```

2. Create a data store using an Oracle database, CustomerDB, by using the following URL:

POST <http://edqserver:8001/edq/config/datasources?pid=4>

An example JSON code to create the data store is:

```
{
  "client":false,
  "name":"CustomersDB",
  "properties":[
    {
      "name":"service",
      "value":"sid"
    },
    {
      "name":"sid",
      "value":"orcl"
    },
    {
      "name":"user",
      "value":"CRM"
    },
    {
      "name":"port",
      "value":"1521"
    },
    {
      "name":"password",
      "value":"welcome123"
    },
    {
      "name":"host",
      "value":"localhost"
    }
  ],
  "species":"oracle"
}
```

 **Note:**

To determine the `pid` or the project ID for the project "Profile Customer", use the HTTP GET operation with the URL:

GET <http://edqserver:8001/edq/config/projects>

3. Create a snapshot by using the following URL:

POST <http://edqserver:8001/edq/snapshots?pid=4>

The JSON code for creating the snapshot is:

```
{
  "name":"CustomersDB.Customers",
  "description":"Customer details",
  "datasource":"CustomersDB",
  "table":"Customers",
  "columns":[
    "ID",
    "FULLNAME",
    "GIVENNAME",
    "FAMILYNAME",
```

```

        "Street",
        "City",
        "State",
        "PostalCode",
        "State",
        "Phone",
        "Cell",
        "Work",
        "eMail",
        "DoB",
        "Gender",
        "Active",
        "CreditLimit",
        "StartDate",
        "EndDate"
    ],
    "sampling":{
        "number":100,
        "offset":0,
        "ordering":"ascending",
        "count":"true"
    }
}

```

The result for this is displayed as:

```

{
  "id": 85,
  "name": "CustomersDB.Customers"
}

```

The snapshot with the name "CustomersDB.Customers" is created.

4. Create a simple process by using the following URL:

POST <http://edqserver:8001/edq/config/processes/simpleprocess?pid=4>

For this example, a simple process is created with a `Quickstats Profiler` and a `Frequency Profiler`, both profiling only the Name fields. This can be done using the following example JSON.

```

{
  "name":"Profile Names",
  "description":"Profile Customer Names",
  "reader":{
    "name":"Read from Customers",
    "stageddata":"Connection to Customers"
  },
  "processors":[
    {
      "name":"Do Quickstats",
      "type":"dn:quickstatsprofiler",
      "columnlist":[
        "GIVENNAME",
        "FAMILYNAME"
      ]
    },
    {
      "name":"Do Frequency Profiling",
      "type":"dn:attributefrequencycountsprofiler"
    }
  ]
}

```

```
    ]  
  }
```

The response to this request is:

```
{  
  "id": 267,  
  "name": "Profile Names"  
}
```

5. Create a simple job by using the following URL:

POST <http://edqserver:8001/edq/jobs/simplejob?pid=4>

```
{  
  "name": "Profile Customer Job",  
  "process": "Profile Names",  
  "description": "Profiling Customer Names",  
  "resultsdrilldown": "none"  
}
```

The response to the request is:

```
{  
  "id": 211,  
  "name": "Profile Customer Job"  
}
```

6. Run the job using the following URL:

POST <http://edqserver:8001/edq/jobs/run>

The JSON code for running the job "Profile Customer Job", which in turn would run the profiling process, is:

```
{  
  "project": "Profile Customer",  
  "job": "Profile Customer Job"  
}
```

The response is:

```
{  
  "executionID": 20,  
  "runeverywhere": false  
}
```

Once this job is running, you can check the status of this execution of the job using the following URL:

GET <http://edqserver:8001/edq/jobs/status?xid=20>

Running this URL displays the status of the job, as shown in the following code:

```
{  
  "complete": true,  
  "endtime": "2016-04-29T14:05:41+01:00",  
  "executionid": 1,  
  "job": "Profile Customer Job",  
  "project": "Profile Customer",  
  "server": "edq_server1",  
  "starttime": "2016-04-29T14:05:38+01:00",  
  "status": "finished"  
}
```

If required, you can cancel the job using the following URL:

```
POST http://edqserver:8001/edq/jobs/cancel
```

The JSON code to cancel a job is:

```
{  
  "executionID": 12345,  
  "type" : "immediate"  
}
```

This would cancel the job instantly, without saving the results. For other options that can be used with "type", see [Cancelling a Running Job](#).

To log in to EDQ Director to view the results of a profiling job that was executed successfully, use the following URL:

```
http://edqserver:8001/edq/blueprints/director/jnlp?  
projectid=1&processid=1&processornum=2
```

The `projectid` and `processid` are the same that are generated using the corresponding REST API calls and the `processornum` value is set to 2, which is the first processor after the reader.

This URL opens the Director UI with the focus on the first profiling processor in the job so that its results can be viewed immediately.

An external application may include an option to remove generated jobs, which would execute calls to the relevant deletion calls. The simplest version of this is to delete the whole project. For details on deleting a project, see [Deleting a Project](#).

10

Using REST APIs for Importing and Exporting Configuration Objects

EDQ 12.2.1.4.4 and later provides a set of REST-based interfaces to automate transfer of configuration between EDQ environments. You can use these REST APIs to package, export, and import the following configuration objects:

- Projects
- Global Reference Data
- Global Data Stores
- Global Published Processors
- Stored Credentials
- Case Sources
- Case Workflows
- Case Permissions

The package file that is written and read by the APIs is in ZIP format. You can encrypt the configuration objects and stored credentials using a password supplied in the request payload.

This chapter provides a detailed description of these interfaces and the operations that can be performed using these interfaces. It includes the following topics:

- [REST Interface for Packaging Configuration](#)
- [JSON Payload Format](#)
- [Packaging Task Status Result Format](#)
- [Packaging REST API Triggers](#)

10.1 REST Interface for Packaging Configuration

The REST interface for working with EDQ packaging, exporting, and importing configuration is

`http://edqserver:port/edq/package`

This interface allows you to perform the following tasks:

- **Exporting configuration**
You must have the **Package** functional permission to package and export configuration objects. If case management objects are included in the export, you must have access to the Case Management Administration application.

To package and export configuration objects, use the following interface:

POST `http://edqserver:port/edq/package/export`

- **Importing configuration**

You must have create and delete function permissions for all the objects that are being imported. If case management objects are included in the import, you must have access to the Case Management Administration application.

To import configuration objects, use the following interface:

```
POST http://edqserver:port/edq/package/import
```

- **Getting packaging task status**

Export and import calls return immediately and the packaging task runs asynchronously. The result of package calls is a JSON object containing an "id" attribute, which is the "execution ID" for the task. For example:

```
{"id": "58636dd0-22b6-4d7d-be16-74908d30404f"}
```

To get status information on a packaging task, use the following interface:

```
GET http://edqserver:port/edq/package/status/executionid
```

Packaging status is retained in the system for one hour after the task has completed. If you do not specify the execution ID in the get status call, the status for all retained packaging tasks is returned.

For example:

```
GET http://server/edq/package/status/58636dd0-22b6-4d7d-be16-74908d30404f
```

returns status information specific to this execution ID.

```
GET http://server/edq/package/status
```

returns status information for all retained packaging tasks. See [Packaging Task Status Result Format](#) for more details.

10.2 JSON Payload Format

The REST APIs use a JSON payload to define the configuration objects for the packaging, export, and import.

The following table lists the attributes of the JSON payload:

Attribute	Description
casemanagement	Defines the case management objects to export/import. See casemanagement attributes .
destination	Required. Destination definition. See destination attributes .
password	Encryption password. If present, configuration objects and stored credentials are encrypted in the package file.
preserve	If true , existing object are not overwritten. Applies to import only.
label	Label describing packaging task. This information is shown in log messages and included in status reports.
datastores	Selector for global data stores to export/import.
publishedprocessors	Selector for global published processors to export/import.
referencedata	Selector for global reference data to export/import.
projects	Selector for projects to export/import.
storedcredentials	Selector for stored credentials to export/import.

casemanagement attributes

The **casemanagement** value is an object with these attributes:

Attribute	Description
sources	Selector for case sources to export/import.
workflows	Selector for case workflows to export/import.
permissions	Selector for case permissions to export/import. Case permissions are matched using the unique internal permission <i>keys</i> , which can be examined in the Case Management Administration application.

destination attributes

The payload **destination** value is an object containing these attributes:

Attribute	Description
location	Required. File name or URL for the package file. If a non-absolute file name is specified, the location is relative to the "local home" configuration directory.
credentials	Name of stored credentials used to authenticate a request to a URL location. Ignored if the location is a file. If the URL requires simple username/password authentication, use "basic" stored credentials.
platformcredentials	If true , uses platform credentials for the file upload/download. Supported on Oracle Cloud Infrastructure (OCI) Object Storage, Amazon Web Services (AWS), and Google Cloud Platform (GCP).
method	HTTP method used for uploads. Ignored if the location is a file. The default is PUT .

Examples:

Use a file in the EDQ landing area:

```
{ ...
  "destination": {
    "location": "landingarea/pkg/package2.zip"
  },
  ...
}
```

Use OCI object storage, with stored credentials:

```
{ ...
  "destination": {
    "location": "https://objectstorage.us-phoenix-1.oraclecloud.com/n/
mytenancy/b/bucket1/o/package.zip",
    "credentials": "OCI 1"
  },
  ...
}
```

Use AWS S3, with stored credentials:

```
{ ...
  "destination": {
    "location": "https://mystorage.s3.us-east-2.amazonaws.com/
package.zip",
    "credentials": "aws1"
  },
  ...
}
```

Use GCP, with stored credentials:

```
{ ...
  "destination": {
    "location": "https://storage.googleapis.com/upload/storage/v1/b/
edqstorage/o?name=package.zip",
    "method": "POST",
    "credentials": "GCP 1"
  },
  ...
}
```

Output Selectors

The selector object used to specify objects to export/import has these attributes:

Attribute	Description
include	Array of "glob-style" patterns used to select objects to include. Use an asterisk (*) to match any characters and a question mark (?) to match a single character.
exclude	Array of "glob-style" patterns used to select objects to exclude. Use an asterisk (*) to match any characters and a question mark (?) to match a single character.
renames	Object specifying the old to new object renames. On export the rename controls the item written to the package. On import the rename applies to objects in the package and controls the name imported into the system.

In an export task the include and exclude patterns are matched against items in the database. In an import task the patterns are matched against the items in the package file.

Examples:

To include all projects except "Temp1" and those starting with "Test", and to include all non-standard reference data:

```
{ ...
  "projects": {
    "include": ["*"],
    "exclude": ["Temp1", "Test*"]
  },
}
```

```
"referencedata": {  
  "include": ["*"],  
  "exclude": ["\\\\"*"]  
},  
...  
}
```

Note the use of the escape character backslash (\) to prevent the first asterisk (*) from being treated as a wild card. To conform to JSON syntax, you need to use two backslash (\) characters.

To include everything:

```
{ ...  
  
  "projects": {  
    "include": ["*"]  
  },  
  
  "referencedata": {  
    "include": ["*"]  
  },  
  
  "datastores": {  
    "include": ["*"]  
  },  
  
  "publishedprocessors": {  
    "include": ["*"]  
  },  
  
  "storedcredentials": {  
    "include": ["*"]  
  },  
  
  "casemanagement": {  
  
    "sources": {  
      "include": ["*"]  
    },  
  
    "workflows": {  
      "include": ["*"]  
    },  
  
    "permissions": {  
      "include": ["*"]  
    }  
  }  
}
```

10.3 Packaging Task Status Result Format

The result of the GET status call is an object in which the attributes are the currently retained execution IDs. The value of each attribute is a status object with these attributes:

Attribute	Description
label	The label copied from the packaging request payload.
type	Whether "import" or "export".
complete	Set to true when the packaging process has completed.
failed	Set to true if the packaging process failed with an errors.
start	Packaging process start timestamp.
end	Packaging process end timestamp. Set when the process is complete.
status	Current status of packaging process. Cleared when process is complete.
error	Error message if failed is true .
manifest	Manifest object defining contents of the export package or the items imported. Set on successful completion.

The manifest for an export task is stored as the first entry in the package ZIP file. The **manifest** object contains these attributes:

Attribute	Description
version	Manifest object version. Always 1 in this version.
appversion	Application version, such as "12.2.1.4.4".
encrypted	Set to true if the package was generated with a password.
projects	Array of names of projects included in the export or matched during import.
referencedata	Array of names of global reference data included in the export or matched during import.
datastores	Array of names of global data stores included in the export or matched during import.
publishedprocessors	Array of names of global published processors included in the export or matched during import.
storedcredentials	Array of names of stored credentials included in the export or matched during import.
casemanagement	Case management items included in the export or matched during import. sources - Case source names workflows - Case workflow names permissions - Case permission keys

Example status result

An export that is run with this payload:

```
{ ...
{ "label": "Export task1",
  "password": "medusa",
  "destination": {
    "location": "https://objectstorage.us-phoenix-1.oraclecloud.com/n/
devbigdata/b/rde.bucket1/o/example.zip",
    "credentials": "OCI 1"
  },
  "projects": {
    "include": ["adb", "cases", "s*"],
    "exclude": ["snowflake"],
    "renames": {
      "adb": "Oracle ADB"
    }
  },
  "referencedata": {
    "include": ["*Country*"]
  },
  "storedcredentials": {
    "include": ["OCI*", "aws*"],
    "exclude": ["aws?"]
  },
  "casemanagement": {
    "sources": {
      "include": ["CS2", "Issue Remediation"]
    },
    "workflows": {
      "include": ["Issue Remediation*"]
    },
    "permissions": {
      "include": ["Permission?"]
    }
  }
}
```

would return this status after completion:

```
{ "d154aee7-9af4-46ab-af89-84ed6b12109a": {
  "start": "2023-07-13T16:44:17.809Z",
  "label": "Export task1",
  "type": "export",
  "end": "2023-07-13T16:44:26.871Z",
```

```
"manifest": {
  "appversion": "14.1.2.0.0",
  "encrypted": true,
  "projects": [
    "Oracle ADB",
    "cases",
    "scannerbug",
    "scripts",
    "services",
    "sqlserver",
    "svr"
  ],
  "referencedata": [
    "Country from City",
    "Nationality to Standard Country",
    "Standardize Country Names"
  ],
  "datastores": [],
  "publishedprocessors": [],
  "storedcredentials": [
    "OCI 1",
    "OCI edqtest",
    "OCI edqtest2",
    "aws - oracle",
    "aws rde",
    "aws s3"
  ],
  "casemanagement": {
    "permissions": [
      "Permission1",
      "Permission2"
    ],
    "workflows": [
      "Issue Remediation Alerts",
      "Issue Remediation Cases"
    ],
    "sources": [
      "CS2",
      "Issue Remediation"
    ]
  },
  "version": 1
},
"complete": true
}
```

10.4 Packaging REST API Triggers

The packaging APIs run triggers with these paths:

- /package/export/start
- /package/export/end

- /package/import/start
- /package/import/end

The arguments to each trigger call are the task label (from the export/import payload) and the packaging status JSON as a string. The status is passed as a string so that it can be sent easily to a logging or streaming framework without additional parsing.

The following is an example that notifies administrators using a web push when an export is complete:

```
addLibrary("webpush")

function getPath() {
  return "/package/export/end"
}

function run(path, id, env, label, status) {
  if (path.endsWith("/end")) {
    var push = WebPush.create("Export " + label + " complete")

    push.title      = "Export notification"
    push.icon       = "images/logo.png"

    push.groupnames = ["Administrators"]
    push.push()
  }
}
```

11

Using REST APIs for User Management and Launchpad Administration

EDQ 12.2.1.4.4 and later includes a set of REST-based interfaces to manage groups, permissions, external group mappings, users, and launchpad applications.

The result for all POST and DELETE calls is an object containing these attributes:

Attribute	Description
ok	Success flag. Possible values are <code>true</code> or <code>false</code> .
message	Error message, present if <code>ok</code> is <code>false</code> .

This chapter includes the following topics:

- [REST Interface for Creating and Updating Groups](#)
- [REST Interface for External Group Mappings](#)
- [REST Interface for Creating and Updating Users](#)
- [REST Interface for Launchpad Applications](#)

11.1 REST Interface for Creating and Updating Groups

The calling user must have the **Access User Administration** functional permission for all user administration calls.

The group object used in requests and responses has these attributes:

Attribute	Type	Description
name	String	The name of the group. Always required.
permissions	Array of strings	Functional permissions associated with group. Values are the internal identifiers representing permissions. In create or update requests, all the permissions in a category can be selected using <i>prefix</i> :. For example ops:* or user_admin:* .
applications	Array of strings	Applications available to the group. Values are the internal identifiers for launchpad applications.

You can perform the following tasks:

- **Get all groups**
To return an array of group objects, use the following interface:

```
GET http://server:port/edq/useradmin/groups
```
- **Get named group**
To get information on a single named group, use the following interface:

GET http://server:port/edq/useradmin/group/name

The result is a group object. If the group does not exist, a 404 response is returned.

- **Create or update single group**

To create or update a single group, use the following interface:

POST http://server:port/edq/useradmin/group

The payload is a single group object. If the `permissions` or `applications` attributes are not specified, the corresponding values in the group are not changed on update. If the group does not exist, the calling user must have the **Add Group** permission, otherwise the user must have the **Modify Group** permission.

- **Create or update multiple groups**

To create or update multiple groups, use the following interface:

POST http://server:port/edq/useradmin/groups

The payload is an array of group objects.

- **Delete named group**

To delete a named group, use the following interface:

DELETE http://server:port/edq/useradmin/group/name

The user must have the **Delete Group** permission.

- **Update group**

To update a group, use the following interface:

POST http://server:port/edq/useradmin/updategroup

This request can be used to add or remove permissions or applications from an existing group. The payload is a JSON object with these attributes:

Attribute	Description
name	The name of the group. Always required.
permissions	Permission additions and removals.
applications	Application additions and removals.

The `permissions` and `applications` attributes are object with these attributes:

Attribute	Description
add	Array of values to add to the group.
remove	Array of values to remove from the group.

Example:

```
{ "name": "Test Group",
  "permissions": {
    "add": [ "user_admin:*", "ops:dnviewallevents"],
    "remove": [ "server_admin:accessserveradmin"]
  },
  "applications": {
    "add": ["opsui" ]
  }
}
```

```
}  
}
```

- **Get lists of supported applications and functional permissions**

Use the following interface:

```
GET http://server:port/edq/useradmin/permissionsinfo
```

Group objects list applications and functional permissions using internal identifiers. The call returns lists of all supported applications and permissions including both the display strings used in the administration web UI and the associated identifiers. Permissions are grouped by the categories shown in the group administration web UI.

The result object contains these attributes:

Attribute	Description
applications	List of supported applications.
functionalpermissions	List of permission categories. Each element contains these attributes: <ul style="list-style-type: none"> – categorylabel - Display string for the category. For example "CM.Static" or "Director". – prefix - Prefix associated with the category. For example "data:" or "user_admin:". – permissions - List of permissions in the category.

Individual application and permission objects contain these attributes:

Attribute	Description
label	Display string for application or permission.
identifier	Internal identifier.

Display strings are returned in the language implied by the web request.

11.2 REST Interface for External Group Mappings

Groups in external realms (LDAP etc) are mapped to EDQ internal group so that the permissions and applications available to a user can be determined on login. Calls in this category return the currently defined mappings and update mappings.

Each mapping is represented by an object with these attributes:

Attribute	Description
realm	The realm name as defined in <code>login.properties</code> . In an update call this attribute may be omitted if exactly one external realm is defined.
externalgroup	The name of group in the external realm. Always required.
internalgroups	Array of EDQ internal group names. In an update call, existing mappings are removed if this attribute is omitted or empty.

You can perform the following tasks:

- **Get current mappings**

To retrieve information about the current external group mappings, use the following interface:

```
GET http://server:port/edq/useradmin/externalgroups
```

The result is an array of mapping objects.

- **Update mappings**

To update information about the current external group mappings, use the following interface:

```
POST http://server:port/edq/useradmin/externalgroups
```

The payload is an array of mapping objects. The calling user must have the **Modify External Group Permissions** functional permission.

Example:

```
[
  { "realm": "EXAMPLE.COM",
    "externalgroup": "edqgroup1",
    "internalgroups": ["Data Analysts"]
  },
  { "realm": "EXAMPLE.COM",
    "externalgroup": "edqgroup2",
    "internalgroups": ["Match Reviewers", "Executives"]
  }
]
```

11.3 REST Interface for Creating and Updating Users

The user object used in requests and responses has these attributes:

Attribute	Type	Description
username	String	The name of the user. Always required.
password	String	The user password. Required when creating a new user.
blocked	Boolean	Set to <code>true</code> if the user has been blocked permanently.
fullname	String	The full name of user. Required when creating a new user.
email	String	The e-mail address of user.
organization	String	The name of the user's organization.
telephone	String	The user's telephone number.
forcepasswordchange	Boolean	If <code>true</code> forces the user to change their password on the next login.
passwordpolicy	Integer	User password expiration policy. Possible values are: <ul style="list-style-type: none"> • 0 - Password expiry defined by system configuration • 1 - Password never expires • 2 - Password expires after system defined period

Attribute	Type	Description
lockoutpolicy	Integer	User lockout policy after bad logins. Possible values are: <ul style="list-style-type: none"> • 0 - Policy defined by system configuration • 1 - No action • 2 - Lockout for 2 minutes • 3 - Lockout for 5 minutes • 4 - Lockout for 15 minutes • 5 - Lockout for 30 minutes • 6 - Lockout for 1 hour • 7 - Lockout for 24 hours • 8 - Lockout permanently
groups	Array of strings	Groups to associate with the user.

The `passwordpolicy` and `lockoutpolicy` attributes correspond to the Password Policy and Invalid Attempts fields in the user definition web UI:

You can perform the following tasks:

- **Get all users**

To return an array of user objects, use the following interface:

```
GET http://server:port/edq/useradmin/users
```

- **Get named user**

To get information on a single named user, use the following interface:

```
GET http://server:port/edq/useradmin/username
```

The result is a user object. If the user does not exist a 404 response is returned.

- **Create or update single user**

To create or update a single user, use the following interface:

```
POST http://server:port/edq/useradmin/group
```

The payload is a single user object. On update, attributes that are not included in the payload object are not modified in the target user. An attribute which is set to **null** in the payload is cleared in the target user object.

For example to update a user and clear the e-mail address, use this payload:

```
{ "username": "user23",
  "email":    null
}
```

To update the password for a user, but force a change on next logon, use the following:

```
{ "username": "user153",
  "password": "newtemppassword",
  "forcepasswordchange" : true
}
```

To create a new user, the caller must have the **Add User** permission. To modify a user, the required permissions depend on the detailed changes. The following table lists the permissions you need corresponding to the information you want to update:

Action	Required Permission
Change password	Change/Reset User Passwords
Block user permanently	Block User
Unblock user	Unblock User
Change full name, e-mail address, organization name or telephone number	Modify User Details
Change password policy, lockout policy or force change option	Modify Account Security Options
Change group membership	Modify User Group Permissions

- **Create or update multiple users**

To create or update multiple users, use the following interface:

```
POST http://server:port/edq/useradmin/users
```

The payload is an array of user objects. To modify users, the required permissions depend on the detailed changes as listed in the table above.

- **Delete named user**

To delete a user, the user must have the **Delete User** permission. Use the following interface:

```
DELETE http://server:port/edq/useradmin/user/username
```

11.4 REST Interface for Launchpad Applications

Calls in this category are used to list and update the applications that are shown on the launchpad for all users. The calling user must have the **Access Server Administration** and **Set User Application Access** permissions for both calls.

You can perform the following tasks:

- **Get applications**

To return the current list of published applications, and a list of all applications known to the system, use the following interface:

```
GET http://server:port/edq/admin/web/launchpad/applications
```

The result object contains these attributes:

Attribute	Type	Description
published	Array of strings	List of applications that are currently shown on the launchpad, in display order.
allapplications	Array of application objects	List of all the applications which are known to the system. Each object contains the following values: <ul style="list-style-type: none"> – label - Display string for application – identifier - Internal identifier <p>Note that the <code>allapplications</code> list includes applications that are permission controlled by group membership (such as Case Management) and also applications that are links to simple content, such as Watchlist Screening Help.</p>

- **Update launchpad applications**

To update launchpad applications, use the following interface:

POST `http://server:port/edq/admin/web/launchpad/applications`

The payload is an array of strings listing the identifiers of the applications that should appear on the launchpad, in order.

For example:

```
[ "opsui",  
  "director",  
  "casemanager",  
  "casemanageradmin"  
]
```

12

Using the Java Messaging Service (JMS) with EDQ

This document describes how to get started using Java Messaging Service (JMS) technology with Oracle Enterprise Data Quality (EDQ). This documentation is intended for system administrators responsible for installing and maintaining EDQ applications.

This chapter includes the following sections:

- [Understanding the Message Queue Architecture](#)
- [Uses of JMS with EDQ](#)
- [Configuring EDQ to Read and Write JMS Messages](#)
- [Defining the Interface Files](#)
- [Illustrations](#)

12.1 Understanding the Message Queue Architecture

JMS is a flexible technology that allows EDQ to integrate with a variety of different messaging systems and technologies, including WebLogic Server Messaging, WebSphere MQ, Oracle Advanced Queueing (OAQ), and many more.

The EDQ server acts as a 'client' of a message queue. The server component may be installed on the same physical server as EDQ, but more commonly the server (or servers) will be remote to it. EDQ is shipped with the client jars needed to connect to either ActiveMQ or Artemis message queue providers. If EDQ is deployed on WebLogic Server, it will also be able to connect natively to WebLogic JMS. Otherwise, for example if EDQ needs to act as a client to an alternative message queue (MQ) system, such as WebSphere MQ, it will need additional client jars to be installed. In this case, the administrator of the MQ system should be consulted to determine which files are needed for a Java application such as EDQ to act as a client.

12.2 Uses of JMS with EDQ

There are two main uses of JMS with EDQ, as listed below: We will focus here on the first use of JMS with EDQ.

- **Consuming and Providing Messages:** You can configure EDQ to read messages from JMS queues, and write messages to them. This can be beneficial where several EDQ servers need to read from a single stream of messages that needs to be persistent to ensure no loss of messages. In this mode, each EDQ server will consume messages from a queue and only 'commit' when it has finished processing each message. Note that JMS is best used for Asynchronous communication, where EDQ is not expected to return a response to a calling application for a specific message.
- **Using JMS in Triggers:** EDQ may send JMS messages to other systems, or may consume JMS messages to start triggers (for example to use a JMS queue to distribute

batch jobs amongst several EDQ servers;). For more details, see the Using JMS in Triggers.

12.3 Configuring EDQ to Read and Write JMS Messages

JMS interfaces to and from EDQ are configured using XML interface files that define:

- The path to the queue of messages
- Properties that define how to work with the specific JMS technology
- How to decode the message payload into a format understood by an EDQ process (for Message Providers – where EDQ reads messages from a queue), or convert messages to a format expected by an external process (for Message Consumers – where EDQ writes messages to a queue).

The XML files are located in the EDQ Local Home directory (formerly known as the config directory), in the following paths:

- buckets/realtime/providers (for interfaces 'in' to EDQ)
- buckets/realtime/consumers (for interfaces 'out' from EDQ)

Once the XML files have been configured, Message Provider interfaces are available in Reader processors in EDQ and to map to Data Interfaces as 'inputs' to a process, and Message Consumer interfaces are available in Writer processors, and to map from Data Interfaces as 'outputs' from a process as shown in Figures 1 and 2 below:

12.4 Defining the Interface Files

An interface file in EDQ consists of three sections, as follows:

- The `<attributes>` section, defining the shape of the interface as understood by EDQ
- The `<messengerconfig>` section, defining how to connect to the JMS queue or topic
- The `<messagebody>` section, defining how to extract contents of a message (e.g. from XML) into attribute data readable by EDQ (for inbound interfaces), or how to convert attribute data from EDQ to message data (e.g. in XML).

12.4.1 Knowing the `<attributes>` section

The `<attributes>` section defines the shape of the interface. It constitutes the attributes that are available in EDQ when configuring a Reader or Writer. For example the following excerpt from the beginning of an interface file configures three string attributes that can be used in EDQ, and their names are:

```
<?xml version="1.0" encoding="UTF-8"?>
<realtimedata messenger="jms">
  <attributes>
<attribute type="string" name="messageID"/>
<attribute type="string" name="name"/>
<attribute type="string" name="AccountNumber"/>
```

```
</attributes>  
  
[file continues]...
```

EDQ supports all the standard attribute types and they are:

- string
- number
- date
- stringarray
- numberarray
- datearray

12.4.2 Knowing the <messengerconfig> Section

The <messengerconfig> section of the interface file configures the settings needed to connect to a given JMS queue or topic on a particular type of JMS technology. The text in <messengerconfig> is parsed as a Java properties file and is merged with the set of properties from the `realtime.properties` file in the EDQ configuration path, with settings in the <messengerconfig> section overriding any matching settings in `realtime.properties`. The `realtime.properties` file therefore allows the configuration of a number of global properties that do not need to be stated in every JMS interface file.

Note:

- As with all property files stored in the EDQ Local Home directory, properties are themselves merged with a set of 'base' properties in the EDQ Home directory which should not be modified, with any property stated in a file in the Local Home directory used in preference to any property in the EDQ Home directory. For example the `realtime.properties` file in the EDQ Local Home directory will be merged with the file of the same name in the EDQ Home directory to form the final set of properties used.
- The `realtime.properties` values may also be used to set properties for Web Services in EDQ. A prefix of `jms.` is therefore used (in this file only, not in the <messengerconfig> section of a JMS interface file) to set properties for JMS (`ws.` is used for Web Services).

The messenger configuration properties, specified either in the <messengerconfig> section of a JMS interface file, or inherited from `realtime.properties`, are used:

- To configure a JNDI name store to look up queues
- To specify the name of the queue/topic and connection factory
- To supply authentication information to the MQ server, if required

JNDI Setup Properties

- `java.naming.factory.initial` - It denotes the class name of the JNDI initial context factory, used to bootstrap JNDI naming

- `java.naming.provider.url` - the JNDI URL used to connect to the JNDI server or local store

In some cases, authentication is required to connect to the JNDI service. In this case the following properties must be added:

- `java.naming.security.principal` - JNDI user name
- `java.naming.security.credentials` - JNDI password

These JNDI properties are a standard Java feature and for more details, see [JNDI Documentation](#).

JMS Names

The other properties are:

- `connectionfactory` - the JNDI name of the JMS 'connection factory'. The default for this (if the property is not set) is 'ConnectionFactory' which is correct for several MQ servers.
- `destination` - the JNDI name of the JMS queue or topic. There is no default for this and it is always required.

Authentication

If the MQ server requires authentication when making connections, add the below properties:

```
username: connection user name  
password: connection password
```



Note:

These properties are used to authenticate against the MQ server and are used to connect to JNDI. In rare cases, both sets may be required.

Notes for specific MQ servers

1. ActiveMQ -

To connect to an ActiveMQ server, use a property setting as listed below:

This can be set in the `<messengerconfig>` section to apply to a single interface, or can be set in `realtime.properties` (prefixed with `jms.`) to set the default for all interfaces.

```
java.naming.provider.url=tcp://host:port
```

ActiveMQ servers generally require connection authentication, so the username and password properties will be required.

To use ActiveMQ on an EDQ server installed on WebLogic, the property settings above need to be entered into the `realtime.properties` in the EDQ Local Home directory. (The settings are present, but commented out, in the file in the Home directory, as it is assumed that WebLogic JMS will normally be used where EDQ is installed on WebLogic.)

2. WebLogic JMS -

When running in a full JavaEE application server, such as WebLogic, Glassfish or WebSphere, the JNDI settings for local JMS are configured automatically and the

JNDI factory and url information do not need to be specified in any EDQ configuration file.

The recommended approach in these cases is to define the JMS configuration within the application server and then just specify the JNDI connection factory and destination names in `<messengerconfig>`.

WebLogic supports the concept of 'foreign' JMS servers. In this case, you define the connection information in the WebLogic Console and expose the destination(s) and connection factory as names in the native WebLogic JNDI store. This works well with Oracle Advanced Queueing (AQ). See below for a snippet of the configuration of a 'foreign JMS server' pointing at an AQ schema.

When JMS is configured like this, a typical `<messengerconfig>` section might be:

```
<messengerconfig>
  connectionfactory = jms/cf1
  destination      = jms/queue1
</messengerconfig>
```

Here `jms/cf1` and `jms/queue1` are JNDI names defined in WebLogic.

12.4.3 Knowing the `<messagebody>` section

This section uses JavaScript to parse message payloads into attributes that EDQ can use for inbound interfaces, and perform the reverse operation (convert EDQ attribute data into message payload data) for outbound interfaces. A function named 'extract' is used to extract data from XML into attribute data for inbound interfaces, and a function named 'build' is used to build XML data from attribute data.

For more details, refer to [Illustrations](#), wherein the scripts are best illustrated using examples.

12.5 Illustrations

Example 1 – Simple Provider File

The following XML is a simple example of a provider interface file, reading messages from a queue in the path `'dynamicQueues/InputQueue'`.

```
<?xml version="1.0" encoding="UTF-8"?>
<realtimedata messenger="jms">
  <attributes>
    <attribute type="string" name="messageID"/>
    <attribute type="string" name="name"/>
    <attribute type="string" name="AccountNumber"/>
    <attribute type="string" name="AccountName"/>
    <attribute type="string" name="Country"/>
  </attributes>
<messengerconfig destination = dynamicQueues/InputQueue </messengerconfig>
  <incoming>
<messagebody>
<script>
<![CDATA[ function extract(str) { var screeningRequest = new XML(str); var
rec = new Record(); rec.messageID = screeningRequest.individual.messageID;
rec.name = screeningRequest.individual.name; rec.AccountNumber =
```

```

screeningRequest.individual.AccountNumber; rec.AccountName =
screeningRequest.individual.AccountName; rec.Country =
screeningRequest.individual.Country; return [rec]; } ]]>
</script>
</messagebody>
<eof>
<messageheader name="JMSType" value="tleof"/>
</eof>
</incoming>
</realtimedata>

```

Example 2 – Simple Consumer File

The following XML is a simple example of a consumer file representing similar data to the provide file listed above, but with additional attributes, which can be added by an EDQ process:

```

<?xml version="1.0" encoding="UTF-8"?>
<realtimedata messenger="jms">
  <attributes>
    <attribute type="string" name="messageID"/>
    <attribute type="string" name="AccountNumber"/>
    <attribute type="string" name="AccountName"/>
    <attribute type="string" name="Country"/>
    <attribute type="string" name="AccountType"/>
  </attributes>
  <messengerconfig> destination = dynamicQueues/OutputQueue </
messengerconfig>
  <outgoing>
    <messagebody>
      <script>
        <![CDATA[ function build(recs) { var rec = recs[0]; var xml =
        <Request> <individual> <messageID>{checkNull(rec.messageID)}</
messageID> <AccountNumber>{checkNull(rec.AccountNumber)}</
AccountNumber> <AccountName>{checkNull(rec.AccountName)}</AccountName>
<Country>{checkNull(rec.Country)}</Country>
<AccountType>{checkNull(rec.AccountType)}</AccountType> </individual>
</Request>; return xml.toXMLString(); } function checkNull(value)
{ return value != null ? value : ""; } ]]>
      </script>
    </messagebody>
  </outgoing>
</realtimedata>

```

Example 3 – Date Parsing and Formatting

The following example file shows how to define a DATE attribute type, and include date parsing and formatting when decoding the message payload:

```

<?xml version="1.0" encoding="UTF-8"?>
<realtimedata messenger="jms">
  <attributes>
    <attribute type="date" name="processingDate"/>
  </attributes>

```

```
<messengerconfig> destination = dynamicQueues/HoldQueue </messengerconfig>
<incoming>
<messagebody>
  <script>
<![CDATA[ var df = Formatter.newDateFormatter("yyyy-MM-dd'T'HH:mm:ss.SSSZ");
function extract(str) { var screeningRequest = new XML(str); var rec = new
Record();rec.processingDate = date(screeningRequest.processingDate; return
[rec]; } function date(x) { var s = x.text().toString(); return s == "" ?
null : df.parse(s); } ]]>
  </script>
</messagebody>
<eof>
<messageheader name="JMSType" value="tleof"/>
</eof>
</incoming>
</realtimeData>
```

13

Using Apache Kafka with EDQ

This document describes how to get started using the Apache Kafka streaming platform with Oracle Enterprise Data Quality (EDQ). This documentation is intended for system administrators responsible for installing and maintaining EDQ applications. In-depth understanding of Kafka concepts and configuration is assumed.



Note:

This feature is applicable only for EDQ 12.2.1.4.1 release.

This chapter includes the following sections:

- [Introduction to Kafka and EDQ](#)
- [Configuring EDQ to Read and Write Kafka Records](#)
- [Defining the Interface Files](#)
- [Illustrations](#)

13.1 Introduction to Kafka and EDQ

Apache Kafka is a highly performant distributed streaming platform.

EDQ can use the [Kafka Consumer API](#) to subscribe to one or more topics and process records as they are published, and can use the [Kafka Producer API](#) to publish a stream of records to a topic.

Kafka records contain a value and an optional key. EDQ supports only text values for record values and keys.

13.2 Configuring EDQ to Read and Write Kafka Records

Kafka interfaces to and from EDQ are configured using XML interface files that define:

- The EDQ attributes produced or consumed by the interface
- Properties that define how to configure the Kafka consumer or producer API
- How to decode the record value into EDQ attributes (for Message Providers – where EDQ reads records from a topic), or convert attributes to a value for a Kafka record (for Message Consumers – where EDQ writes messages to a topic).

The XML files are located in the EDQ Local Home directory, in the following paths:

- buckets/realtime/providers (for interfaces 'in' to EDQ)
- buckets/realtime/consumers (for interfaces 'out' from EDQ)

Once the XML files have been configured, Message Provider interfaces are available in Reader processors in EDQ and to map to Data Interfaces as 'inputs' to a process, and Message Consumer interfaces are available in Writer processors and to map from Data Interfaces as 'outputs' from a process, in the same way as Web Service inputs/ outputs and JMS providers/consumers.

13.3 Defining the Interface Files

An interface file in EDQ consists of a `realtimedata` element which defines the message framework. For Kafka interfaces, use the following:

```
<realtimedata messenger="kafka">
  ...
</realtimedata>
```

The `realtimedata` element contains three subsections:

- The `<attributes>` section, defining the shape of the interface as understood by EDQ
- The `<messengerconfig>` section, defining how to configure the Kafka API
- A message format section defining how a Kafka record is mapped from or to EDQ attributes. For provider interfaces, the element is `<incoming>`; for consumer interfaces, the element is `<outgoing>`.

13.3.1 Understanding the `<attributes>` section

The `<attributes>` section defines the shape of the interface. It configures the attributes that are available in EDQ when configuring a Reader or Writer. For example, the following excerpt from the beginning of an interface file configures string and number attributes that can be used in EDQ:

```
<?xml version="1.0" encoding="UTF-8"?>
<realtimedata messenger="kafka">
  <attributes>
    <attribute type="string" name="messageID"/>
    <attribute type="string" name="name"/>
    <attribute type="number" name="AccountNumber"/>
  </attributes>
  ...
</realtimedata>
```

EDQ supports all the standard attribute types. These are:

- string
- number
- date
- stringarray
- numberarray
- datearray

13.3.2 Understanding the <messengerconfig> section

The <messengerconfig> section of the interface file configures the Kafka API. The text in <messengerconfig> is parsed as a set of Java properties.

Properties prefixed with `conf.` are passed directly to the Kafka API after removing the `conf.` prefix from the key.

Other properties which may be placed in the <messengerconfig> section for Kafka are:

- `topic`: For provider interfaces, a comma or space separated list of Kafka topics to subscribe to is required. For consumer interfaces, a single topic name is required.
- `poll`: The interval between polls for new records in milliseconds. This is only applicable for provider interfaces. The default value is 500.
- `key.encoding` and `value.encoding`: The character sets used for converting record keys and values. The character sets must be recognized by `java.nio.charset.Charset.forName`. The default values are implementation-specific.

On a WebLogic installation, authentication information required in configuration properties may be stored in the OPSS credentials store, and used in properties with `${username}` and `{password}` substitution. Use the `cred.key` and `cred.map` properties to define the credentials store key and map names. The map name defaults to "edq" if omitted.

The following is an example of a complete configuration with credentials store use:

```
<messengerconfig>
  cred.key                = kafka1
  topic                  = mytopic

  conf.bootstrap.servers = kserver:9094
  confs.acks              = all
  conf.max.block.ms      = 1000
  conf.security.protocol = SASL_SSL
  conf.sasl.mechanism    = PLAIN
  conf.ssl.truststore.location = pathtokeystore.jks
  conf.ssl.truststore.password = pw
  conf.sasl.jaas.config   =
org.apache.kafka.common.security.plain.PlainLoginModule required
  username="${username}" password="{password}";
</messengerconfig>
```

The user name and password for the `jaas.config` property are read from the credentials store with key "kafka1".

Default properties may be defined in the `realtime.properties` file in the EDQ local home directory. Keys for Kafka interfaces in this file are prefixed with "kafka.". For example:

```
kafka.conf.security.protocol = SASL_SSL
```

13.3.3 Understanding the <incoming> or <outgoing> section

The `<incoming>` or `<outgoing>` section defines how record metadata and values are converted to/from EDQ attributes. It consists of the following two subsections:

- [The `<messageheaders>` section](#)
- [The `<messagebody>` section](#)

13.3.3.1 Understanding the `<messageheaders>` section

The `<messageheaders>` section is optional. It defines how data outside the record value is converted to/from EDQ attributes.

The format of this section is as follows:

```
<messageheaders>
  <header name="headername" attribute="attributename"/>
  ...
</messageheaders>
```

The Kafka interface defines two header names as follows:

- **key:** For providers, the key value from a record is stored in the named EDQ attribute. For consumers, the value of the EDQ attribute is used as the record key on publish.
- **topic:** The name of the Kafka topic on which a record was received is stored in the EDQ attribute. It is only applicable for consumers. This is useful if the interface is defined to subscribe to a number of different topics.

13.3.3.2 Understanding the `<messagebody>` section

The `<messagebody>` section defines how the text value in a Kafka message is converted from/to EDQ attributes. The element is followed by a subsection defining the conversion mechanism. The following conversion mechanisms are supported:

- [JSON Conversion](#)
- [Script Conversion](#)

JSON Conversion

The record value is expected to be in JSON format. The nested elements define the mapping between JSON attributes and EDQ attributes.

The format is as follows:

```
<json [multirecord="true or false"] [defaultmappings="true or false"]>
  <mapping attribute="attributename" path="jsonattributename"/>
  ...
</json>
```

If the `defaultmappings` attribute is omitted or set to **true**, automatic mappings are created from the interface EDQ attributes to JSON attributes.

If the `multirecord` attribute is set to **true**, consumer interfaces expect the JSON input to be an array and provider interfaces to generate a JSON array.

Assuming the attributes shown in the `<attributes>` section description, here are some JSON conversion examples:

Example 1:

An example of a simple JSON conversion in which everything is automatic is as follows:

```
<json/>
```

A simple conversion with automatic mappings expects and generates values like:

```
{ "messageID": "x123",  
  "name": "John Smith",  
  "AccountNumber": 34567  
}
```

Example 2:

An example of a multirecord JSON conversion with a mapping is as follows:

```
<json multirecord="true">  
  <mapping attribute="AccountNumber" path="accno"/>  
</json>
```

A multirecord conversion with a single attribute name mapping expects and generates values like:

```
[{ "messageID": "x123",  
  "name": "John Smith",  
  "accno": 34567  
},  
...  
]
```

Script Conversion

To support more complex conversions, for example XML parsing, a JavaScript script can be provided to process the record value.

In provider interfaces, the script must define a function named "extract" which takes the string record value as an argument. The script should return an array of **Record** objects with attribute names matching EDQ attributes.

The following is an example parsing some XML, using the E4X XML processing API in Rhino JavaScript:

```
<script>  
<![CDATA[  
  function extract(str) {  
    var r = new Record()  
    var x = new XML(XMLTransformer.purifyXML(str));
```

```

        r.messageID      = x.ID
        r.name           = x.Accname
        r.accountNumber = parseInt(x.Accnumber)
        return [r];
    }
  ]]>
</script>

```

In consumer interfaces, the script must define a function named "build" which takes an array of **Record** objects and returns the text value.

The following is an example generating some XML:

```

<script>
  <![CDATA[
    function build(recs, mtags) {
      var rec = recs[0];

      var xml =
        <response xmlns="http://www.datanomic.com/ws">
          <sum>{rec.sum}</sum>
        </response>;

      return xml.toXMLString();
    }
  ]]>
</script>

```

In a multi-record response, the default behaviour is to call the script for each record. If `<script multirecord="true">` is used, the build function is called once with all the records in the message.

For more details, refer to [Illustrations](#), which provides an example of a provider interface file using default JSON conversion.

13.4 Illustrations

The following XML is a simple example of a provider interface file, using default JSON conversion:

```

<?xml version="1.0" encoding="UTF-8"?>
<realtimedata messenger="kafka">

  <attributes>
    <attribute type="string" name="messageID"/>
    <attribute type="string" name="name"/>
    <attribute type="string" name="AccountNumber"/>
    <attribute type="string" name="AccountName"/>
    <attribute type="string" name="Country"/>
  </attributes>

  <messengerconfig>

```

```
cred.key                = kafka1
topic                  = mytopic

conf.bootstrap.servers = kserver:9094
confs.acks              = all
conf.max.block.ms       = 1000
conf.security.protocol  = SASL_SSL
conf.sasl.mechanism     = PLAIN
conf.ssl.truststore.location = mykeystore,jks
conf.ssl.truststore.password = pw
conf.sasl.jaas.config   =
org.apache.kafka.common.security.plain.PlainLoginModule required
  username="${username}" password="${password}";
</messengerconfig>

<incoming>
  <messagebody>
    <json/>
  </messagebody>
</incoming>
</realtimedata>
</realtimedata>
```

14

Using Amazon Simple Queue Service (Amazon SQS) with EDQ

This document describes how to get started using the Amazon Simple Queue Service (Amazon SQS) technology with Oracle Enterprise Data Quality (EDQ). This documentation is intended for system administrators responsible for installing and maintaining EDQ applications.

This chapter includes the following sections:

- [Introduction to Amazon SQS and EDQ](#)
- [Configuring EDQ to Read and Write Amazon SQS Messages](#)
- [Defining the Interface Files](#)
- [Illustrations](#)

14.1 Introduction to Amazon SQS and EDQ

EDQ realtime provider and consumer buckets can be configured to use Amazon SQS queues for reading and publishing.

14.2 Configuring EDQ to Read and Write Amazon SQS Messages

Amazon SQS interfaces to and from EDQ are configured using XML interface files that define:

- The path to the queue of messages
- Properties that define how to work with the specific Amazon SQS technology
- How to decode the message payload into a format understood by an EDQ process (for Message Providers – where EDQ reads messages from a queue), or convert messages to a format expected by an external process (for Message Consumers – where EDQ writes messages to a queue).

The XML files are located in the EDQ Local Home directory (formerly known as the config directory), in the following paths:

- buckets/realtime/providers (for interfaces 'in' to EDQ)
- buckets/realtime/consumers (for interfaces 'out' from EDQ)

Once the XML files have been configured, Message Provider interfaces are available in Reader processors in EDQ and to map to Data Interfaces as 'inputs' to a process, and Message Consumer interfaces are available in Writer processors, and to map from Data Interfaces as 'outputs' from a process.

14.3 Defining the Interface Files

An interface file in EDQ consists of a `realtimedata` element which defines the message framework. For Amazon SQS interfaces, use the following:

```
<?xml version="1.0" encoding="UTF-8"?>

<realtimedata messenger="sqs">
    ...
</realtimedata>
```

The `realtimedata` element contains three subsections:

- The `<attributes>` section, defining the shape of the interface as understood by EDQ
- The `<messengerconfig>` section, defining how to connect to the Amazon SQS queue.
- A message format section defining how to extract contents of a message (e.g. from XML) into attribute data readable by EDQ (for inbound interfaces), or how to convert attribute data from EDQ to message data (e.g. in XML). For provider interfaces, the element is `<incoming>`; for consumer interfaces, the element is `<outgoing>`.

14.3.1 Understanding the `<attributes>` section

The `<attributes>` section defines the shape of the interface. It constitutes the attributes that are available in EDQ when configuring a Reader or Writer. For example, the following excerpt from the beginning of an interface file configures string and number attributes that can be used in EDQ:

```
<?xml version="1.0" encoding="UTF-8"?>
<realtimedata messenger="sqs">
<attributes>
    <attribute type="string" name="messageID"/>
    <attribute type="string" name="name"/>
    <attribute type="number" name="AccountNumber"/>
</attributes>
```

[file continues]...

EDQ supports all the standard attribute types and they are:

- string
- number
- date
- stringarray
- numberarray
- datearray

14.3.2 Understanding the <messengerconfig> section

The following properties can be set in the <messengerconfig> section:

Property	Description
queue	SQS queue URL (required).
credentials	Stored credentials name used to connect to stream; if omitted, platform (instance) authentication is used.
interval	Interval in milliseconds between polls for message reception. The default is 0 (long polls are used).
proxy	host:port proxy server for HTTPS calls.
deletemode	Controls message deletion after reception. Valid values are: <ul style="list-style-type: none"> • reception: Each message is deleted immediately on reception. This is the default. • completion: Each message is deleted when it has completed traversing the processes in the job. • off: Messages are not deleted automatically. They must be deleted manually, perhaps using the reception handle in a web service call.
MessageDeduplicationId	Deduplication ID for FIFO queues. May be overridden by header attribute. Two special values are supported: <ul style="list-style-type: none"> • \$hash: The deduplication ID is computed as the SHA-256 hash of the message body. • \$uuid: A random UUID is used as the deduplication ID.
MessageGroupId	Group ID for FIFO queues. May be overridden by header attribute.
DelaySeconds	Sending delay in seconds. May be overridden by header attribute.
MaxNumberOfMessages	Maximum number of messages to receive in one call. The default is 10.
VisibilityTimeout	The duration (in seconds) that the received messages are hidden from subsequent receive requests. The default is set in the queue definition.
WaitTimeSeconds	Wait time in seconds for long poll receive requests. The default is 20s.

Defaults can be set in `realtime.properties` with prefix "sqs."

14.3.3 Understanding the <incoming> or <outgoing> section

The <incoming> or <outgoing> section defines how message metadata and values are converted to/from EDQ attributes. It consists of the following two subsections:

- [The <messageheaders> section](#)

- [The <messagebody> section](#)

14.3.3.1 Understanding the <messageheaders> section

The <messageheaders> section allows attributes to be mapped to additional sending properties for transmission, and attributes to be set from message metadata on reception. Refer to the [SQS Documentation](#) for more details on each value. A non empty message header value will override a <messengerconfig> property with the same name.

The following standard headers are available:

Header name	Settable	Readable	Type	Description
DelaySeconds	yes	no	number	Sending delay in seconds
MessageDeduplicationId	yes	yes	string	Deduplication ID for FIFO queues
MessageGroupId	yes	yes	string	Group ID for FIFO queues
ApproximateFirstReceiveTimestamp	no	yes	number	Timestamp for first receive from queue
ApproximateReceiveCount	no	yes	number	Number of times message has been received
AWSTraceHeader	no	yes	string	X-Ray trace header
SenderId	no	yes	string	IAM user or role ID
SentTimestamp	no	yes	number	Sending timestamp
SequenceNumber	no	yes	number	Sequence number from SQS
MessageId	no	yes	string	Internal message ID
ReceiptHandle	no	yes	string	Receipt handle, required for manual deletion of messages

Custom message attributes

In addition to the standard headers defined above, the <messageheaders> section can also define custom attributes. Custom attribute names are prefixed with "messageattribute:" to distinguish them from standard headers.

Example custom attribute

```
<messageheaders>
  <header name="messageattribute:tel" attribute="telephone"
```

```
type="string"/>
</messageheaders>
```

14.3.3.2 Understanding the <messagebody> section

This section uses JavaScript to parse message payloads into attributes that EDQ can use for inbound interfaces, and perform the reverse operation (convert EDQ attribute data into message payload data) for outbound interfaces. A function named 'extract' is used to extract data from XML into attribute data for inbound interfaces, and a function named 'build' is used to build XML data from attribute data.

For more details, refer to [Illustrations](#), which provides an example of a complete provider bucket which can receive JSON messages from a case management filter reporting trigger.

14.4 Illustrations

The following XML is a simple example of a complete provider bucket which can receive JSON messages from a case management filter reporting trigger. The sender ID and message group ID are also returned, along with two custom message attributes, **attr1** and **attr2**.

```
<?xml version="1.0" encoding="UTF-8"?>

<realtimedata messenger="sqs">
  <attributes>
    <attribute name="filter" type="string"/>
    <attribute name="type" type="string"/>
    <attribute name="xaxis" type="string"/>
    <attribute name="yaxis" type="string"/>
    <attribute name="server" type="string"/>
    <attribute name="userid" type="number"/>
    <attribute name="user" type="string"/>
    <attribute name="userdisplay" type="string"/>
    <attribute name="start" type="date" format="iso"/>
    <attribute name="duration" type="number"/>
    <attribute name="status" type="string"/>
    <attribute name="sql" type="string"/>
    <attribute name="arg.type" type="stringarray"/>
    <attribute name="arg.value" type="stringarray"/>
    <attribute name="senderid" type="string"/>
    <attribute name="attr1" type="string"/>
    <attribute name="attr2" type="number"/>
    <attribute name="mgid" type="string"/>
  </attributes>

  <messengerconfig>
    queue = https://sqs.eu-west-1.amazonaws.com/458503484332/queue1
    credentials = aws1
    deletemode = completion
  </messengerconfig>

  <incoming>

    <messageheaders>
```

```
        <header name="SenderId"           attribute="senderid"/>
        <header name="MessageGroupId"     attribute="mgid"/>
        <header name="messageattribute:attr1" attribute="attr1"
type="string"/>
        <header name="messageattribute:attr3" attribute="attr2"
type="number"/>
    </messageheaders>

    <messagebody>
        <script>
            <![CDATA[
                var simple = ["filter", "type", "xaxis", "yaxis", "server",
"userid", "user", "userdisplay", "duration", "status", "sql"]

                function extract(str) {
                    var obj = JSON.parse(str)
                    var rec = new Record()

                    for (let x of simple) {
                        rec[x] = obj[x]
                    }

                    rec.start = obj.start && new Date(obj.start)

                    if (obj.args) {
                        rec['arg.type'] = obj.args.map(a => a.type)
                        rec['arg.value'] = obj.args.map(a => a.value &&
a.value.toString())
                    }

                    return [rec];
                }
            ]]>
        </script>
    </messagebody>

</incoming>

</realtimedata>
```

15

Using Oracle Cloud Infrastructure (OCI) Queue with EDQ

This document describes how to get started using the Oracle Cloud Infrastructure (OCI) Queue service with Oracle Enterprise Data Quality (EDQ). This documentation is intended for system administrators responsible for installing and maintaining EDQ applications.

This chapter includes the following sections:

- [Introduction to OCI Queue and EDQ](#)
- [Configuring EDQ to Read and Write OCI Queue Messages](#)
- [Defining the Interface Files](#)
- [Illustrations](#)

15.1 Introduction to OCI Queue and EDQ

EDQ realtime provider and consumer buckets can be configured to use OCI Queue service for reading and publishing. This queue service is a much better fit for realtime processing as compared to OCI streams.

15.2 Configuring EDQ to Read and Write OCI Queue Messages

OCI Queue interfaces to and from EDQ are configured using XML interface files that define:

- The path to the queue of messages
- Properties that define how to work with the specific OCI Queue technology
- How to decode the message payload into a format understood by an EDQ process (for Message Providers – where EDQ reads messages from a queue), or convert messages to a format expected by an external process (for Message Consumers – where EDQ writes messages to a queue).

The XML files are located in the EDQ Local Home directory (formerly known as the config directory), in the following paths:

- buckets/realtime/providers (for interfaces 'in' to EDQ)
- buckets/realtime/consumers (for interfaces 'out' from EDQ)

Once the XML files have been configured, Message Provider interfaces are available in Reader processors in EDQ and to map to Data Interfaces as 'inputs' to a process, and Message Consumer interfaces are available in Writer processors, and to map from Data Interfaces as 'outputs' from a process.

15.3 Defining the Interface Files

An interface file in EDQ consists of a `realtimedata` element which defines the message framework. For OCI Queue interfaces, use the following:

```
<?xml version="1.0" encoding="UTF-8"?>
<realtimedata messenger="ociqueues">
  ...
</realtimedata>
```

The `realtimedata` element contains three subsections:

- The `<attributes>` section, defining the shape of the interface as understood by EDQ
- The `<messengerconfig>` section, defining how to connect to OCI Queue.
- A message format section defining how to extract contents of a message (e.g. from XML) into attribute data readable by EDQ (for inbound interfaces), or how to convert attribute data from EDQ to message data (e.g. in XML). For provider interfaces, the element is `<incoming>`; for consumer interfaces, the element is `<outgoing>`.

15.3.1 Understanding the `<attributes>` section

The `<attributes>` section defines the shape of the interface. It constitutes the attributes that are available in EDQ when configuring a Reader or Writer. For example, the following excerpt from the beginning of an interface file configures string and number attributes that can be used in EDQ:

```
<?xml version="1.0" encoding="UTF-8"?>
<realtimedata messenger="ociqueues">
<attributes>
  <attribute type="string" name="messageID"/>
  <attribute type="string" name="name"/>
  <attribute type="number" name="AccountNumber"/>
</attributes>
```

[file continues]...

EDQ supports all the standard attribute types and they are:

- string
- number
- date
- stringarray
- numberarray
- datearray

15.3.2 Understanding the <messengerconfig> section

The following properties can be set in the <messengerconfig> section:

Property	Description
queue	The queue OCID (required).
credentials	Stored credentials name used to connect to stream; if omitted, platform (instance) authentication is used.
interval	Interval in milliseconds between polls for message reception. The default is 0 (long polls are used).
proxy	host:port proxy server for HTTPS calls.
deletemode	Controls message deletion after reception. Valid values are: <ul style="list-style-type: none"> reception: Each message is deleted immediately on reception. This is the default. completion: Each message is deleted when it has completed traversing the processes in the job. off: Messages are not deleted automatically. They must be deleted manually, perhaps using the reception handle in a web service call.
limit	Maximum number of messages to receive in one call. The default is 20.
visibilityInSeconds	The duration (in seconds) that the received messages are hidden from subsequent receive requests. The default is set in the queue definition.
timeoutInSeconds	Wait time in seconds for long poll receive requests. If omitted the OCI default of 30s is used.

Defaults can be set in `realtime.properties` with the prefix "ociqueues."

15.3.3 Understanding the <incoming> or <outgoing> section

The <incoming> or <outgoing> section defines how message metadata and values are converted to/from EDQ attributes. It consists of the following two subsections:

- [The <messageheaders> section](#)
- [The <messagebody> section](#)

15.3.3.1 Understanding the <messageheaders> section

The <messageheaders> section allows attributes to be mapped to additional sending properties for transmission, and attributes to be set from message metadata on reception. Refer to the [OCI Queue Documentation](#) for more details on each value. A non empty message header value will override a <messengerconfig> property with the same name.

The following standard headers are available:

Header name	Settable	Readable	Type	Description
deliveryCount	no	yes	number	The number of times the message has been delivered to a consumer.

Header name	Settable	Readable	Type	Description
expireAfter	no	yes	date	The time after which the message will be automatically deleted.
id	no	yes	string	The internal message ID.
receipt	no	yes	string	The Receipt token that is required for manual deletion of messages.
visibleAfter	no	yes	date	The time after which the message will be visible to other consumers.

15.3.3.2 Understanding the <messagebody> section

This section uses JavaScript to parse message payloads into attributes that EDQ can use for inbound interfaces, and perform the reverse operation (convert EDQ attribute data into message payload data) for outbound interfaces. A function named 'extract' is used to extract data from XML into attribute data for inbound interfaces, and a function named 'build' is used to build XML data from attribute data.

For more details, refer to [Illustrations](#), which provides an example of a complete provider bucket which can receive JSON messages from a case management filter reporting trigger.

15.3.4 Illustrations

The following XML is a simple example of a complete provider bucket which can receive JSON messages from a case management filter reporting trigger. The internal message ID is also returned.

```
<?xml version="1.0" encoding="UTF-8"?>

<realtimedata messenger="ociqueues">
  <attributes>
    <attribute name="filter"      type="string"/>
    <attribute name="type"       type="string"/>
    <attribute name="xaxis"      type="string"/>
    <attribute name="yaxis"      type="string"/>
    <attribute name="server"     type="string"/>
    <attribute name="userid"     type="number"/>
    <attribute name="user"       type="string"/>
    <attribute name="userdisplay" type="string"/>
    <attribute name="start"      type="date" format="iso"/>
    <attribute name="duration"   type="number"/>
    <attribute name="status"     type="string"/>
    <attribute name="sql"        type="string"/>
    <attribute name="arg.type"   type="stringarray"/>
    <attribute name="arg.value"  type="stringarray"/>
    <attribute name="mgid"       type="string"/>
  </attributes>

  <messengerconfig>
    queue =
ocid1.queue.oc1.phx.amaaaaaa7u6obfiaotsz7yuj2zcbc5uhc45evvv7wfwedgertw3
```

```
543we
  credentials = OCI 1
</messengerconfig>

<incoming>

  <messageheaders>
    <header name="id" attribute="mgid"/>
  </messageheaders>

  <messagebody>
    <script>
      <![CDATA[
        var simple = ["filter", "type", "xaxis", "yaxis", "server",
"userid", "user", "userdisplay", "duration", "status", "sql"]

        function extract(str) {
          var obj = JSON.parse(str)
          var rec = new Record()

          for (let x of simple) {
            rec[x] = obj[x]
          }

          rec.start = obj.start && new Date(obj.start)

          if (obj.args) {
            rec['arg.type'] = obj.args.map(a => a.type)
            rec['arg.value'] = obj.args.map(a => a.value &&
a.value.toString())
          }

          return [rec];
        }
      ]]>
    </script>
  </messagebody>

</incoming>

</realtimedata>
```

16

Using Scripted Global Web Services with EDQ

EDQ 12.2.1.4.4 restores limited support for global web services. This documentation is intended for system administrators responsible for installing and maintaining EDQ applications.

This chapter includes the following topics:

- [Introduction to Global Web Services and EDQ](#)
- [Configuring EDQ to Read and Write Web Service Requests](#)
- [Defining the Interface Files](#)
- [Illustrations](#)

16.1 Introduction to Global Web Services and EDQ

EDQ 12.2.1.4.0 and earlier supported "global" SOAP web services, which were installed as jars in the local webservices directory. Support for global web services was removed in EDQ 12.2.1.4.1.

EDQ 12.2.1.4.4 onwards you can configure realtime provider and consumer buckets to use global web services for reading and publishing. These web services bucket files use scripts to decode incoming text payloads and generate outgoing text results. The supported payload formats are JSON (application/json), XML (text/xml), and plain text (text/plain).



Note:

SOAP is not supported.

16.2 Configuring EDQ to Read and Write Web Service Requests

Web service interfaces to and from EDQ are configured using XML interface files that define:

- The URL path for the web service
- How to decode the message payload into a format understood by an EDQ process (for Message Providers – where EDQ reads messages), or convert messages to a format expected by an external process (for Message Consumers – where EDQ writes messages).

The XML files are located in the EDQ Local Home directory (formerly known as the config directory), in the following paths:

- buckets/realtime/providers (for interfaces 'in' to EDQ)
- buckets/realtime/consumers (for interfaces 'out' from EDQ)

Once the XML files have been configured, Message Provider interfaces are available in Reader processors in EDQ and to map to Data Interfaces as 'inputs' to a process, and Message Consumer interfaces are available in Writer processors, and to map from Data Interfaces as 'outputs' from a process.

16.3 Defining the Interface Files

An interface file in EDQ consists of a `realtimedata` element which defines the message framework. A web service that returns a result is defined by provider and consumer buckets. The configuration in both must be marked as synchronous. For web services, use the following:

```
<?xml version="1.0" encoding="UTF-8"?>

<realtimedata messenger="ws" synchronous="true">
  ...
</realtimedata>
```

The `realtimedata` element contains three subsections:

- The `<attributes>` section, defining the shape of the interface as understood by EDQ
- The `<messengerconfig>` section, defining how to connect to the web service endpoint.
- A message format section defining how to extract contents of a message (e.g. from XML) into attribute data readable by EDQ (for inbound interfaces), or how to convert attribute data from EDQ to message data (e.g. in XML). For provider interfaces, the element is `<incoming>`; for consumer interfaces, the element is `<outgoing>`.

16.3.1 Understanding the `<attributes>` section

The `<attributes>` section defines the shape of the interface. It constitutes the attributes that are available in EDQ when configuring a Reader or Writer.

Provider scripts generate `record` object from the incoming payload and consumer scripts convert record contents into the outgoing payload. In a provider bucket script, create a record using `new Record()` or `newRecord()`. The second form should be used in Groovy scripts. Then set the attributes with assignment like:

```
rec.id = 23
rec.name = "John Smith"
```

Record objects contain the attributes defined in the bucket's attributes section. For example, the following excerpt from the beginning of an interface file defines the bucket's **attributes**, which generate record objects that contain **id**, **name** and **email** attributes:

```
<?xml version="1.0" encoding="UTF-8"?>
<realtimedata messenger="ws">
<attributes>
```

```
<attribute name="id" id="1" type="number"/>
<attribute name="name" id="2" type="string"/>
<attribute name="email" id="3" type="string"/>
</attributes>
```

[file continues]...

EDQ supports all the standard attribute types and they are:

- string
- number
- date
- stringarray
- numberarray
- datearray

16.3.2 Understanding the <messengerconfig> section

The `messengerconfig` section must include a path property. The value is appended to `http://server/edq/restws/` to form the endpoint for the web service.

The path can contain letters, digits, and underscore (`_`) characters.

16.3.3 Understanding the <incoming> or <outgoing> section

The `<incoming>` or `<outgoing>` section defines how message metadata and values are converted to/from EDQ attributes. It consists of the following two subsections:

- [Understanding the <messageheaders> section](#)
- [Understanding the <messagebody> section](#)

The provider and consumer buckets are linked by adding a key attribute to the incoming and outgoing elements:

```
<incoming key="addup1">
<outgoing key="addup1">
```

The same key value must be used in the provider and consumer web service buckets for a single service. Each service should use different key values.

16.3.3.1 Understanding the <messageheaders> section

The `<messageheaders>` section is optional. It defines how data outside the record value is converted to/from EDQ attributes. The format of this section is as follows:

```
<messageheaders>
  <header name="headername" attribute="attributename"/>
  ...
</messageheaders>
```

16.3.3.2 Understanding the <messagebody> section

This section uses JavaScript to parse message payloads into attributes that EDQ can use for inbound interfaces, and perform the reverse operation (convert EDQ attribute data into message payload data) for outbound interfaces. A function named 'extract' is used to extract data from XML into attribute data for inbound interfaces, and a function named 'build' is used to build XML data from attribute data.

Provider Definition

The provider bucket XML file uses a script decoder to create attributes from the input payload:

```
<messagebody>
  <script content="...">
    <![CDATA[
      function extract(payload, type, mtags, env) {
        ...
      }
    ]]>
  </script>
</messagebody>
```

The **content** attribute is a comma-separated list of the input types that the provider supports:

- **json**: JSON text with content type *application/json*.
- **dom**: XML text with content type *text/xml*. The XML is parsed to DOM by the web services framework. To decode the DOM in the script use `new XML(payload)`.
- **text**: Plain text with content type *text/plain*.

A web service call with a content type not supported by the script will fail with a 415 error. The actual type used is available in `env.contentType`. If the script supports more than one content type, you can use this to determine how to parse the payload.

For more details, refer to [Illustrations](#), which provides examples of a provider bucket XML file.

Consumer Definition

The consumer bucket XML file uses a script encoder to create the result text from input attributes:

```
<messagebody>
  <script content="...">
    <![CDATA[
      function build(recs, mtags, ctype) {
        ...
      }
    ]]>
  </script>
</messagebody>
```

The **content** attribute is a comma-separated list of the output types that the consumer supports:

- **recs** is an array of record objects.
- **ctype** is the content type (json, dom, or text) of the incoming request. If the service can accept more than one content type, use this to create the output in the correct format.

Set `multirecord="true"` if a request can contain multiple records:

```
<script multirecord="true">
```

For more details, refer to [Illustrations](#), which provides examples of a consumer bucket XML file.

16.4 Illustrations

This is an example of a simple global web service called *addup*, which accepts two numbers and generates the sum. Supported input formats are JSON, XML, and plain text. The path for the service is *addup* and the web service endpoint is:

```
http://server:port/edq/restws/addup
```

The following table illustrates how the input and output formats work based on this example service.

Supported Format	Example Input	Example Output
JSON	<pre>{ "n1" : 100, "n2" : 235 }</pre>	<pre>{ "inputs": [100, 235], "sum": 335 }</pre>
XML	<pre><request> <n1>100</n1> <n2>223</n2> </request></pre>	<pre><response> <inputs> <v>100</v> <v>223</v> </inputs> <sum>323</sum> </response></pre>
Plain text	234,129	234 + 129 = 363

Example 1 – Provider File

The following XML is an example of a provider bucket XML file:

```
<?xml version="1.0" encoding="UTF-8"?>
<realtimedata messenger="ws" synchronous="true">
```

```
<!-- ===== -->
<!-- The attributes which are created by this provider -->
<!-- ===== -->

<attributes>
  <attribute name="n1" id="1" type="number"/>
  <attribute name="n2" id="2" type="number"/>
</attributes>

<!-- ===== -->
<!-- The web service configuration -->
<!-- ===== -->

<messengerconfig>
  path = addup
</messengerconfig>

<!-- ===== -->
<!-- The script which unpacks the payload -->
<!-- ===== -->

<incoming key="addup1">
  <messagebody>
    <script content="json,text,dom">
      <![CDATA[
        function extract(str, type, mtags, env) {
          if (env.contenttype == "json") {
            var x = JSON.parse(str)
            var r = new Record();

            r.n1 = x.n1
            r.n2 = x.n2

            return [r];
          } else if (env.contenttype == "dom") {
            var x = new XML(str)
            var r = new Record();

            r.n1 = parseInt(x.n1)
            r.n2 = parseInt(x.n2)

            return [r];
          } else {
            var a = str.split(",")
            var r = new Record();

            r.n1 = parseInt(a[0])
            r.n2 = parseInt(a[1])

            return [r];
          }
        }
      ]]>
    </script>
```

```

    </messagebody>
  </incoming>

</realtimedata>

```

Example 2 – Consumer File

The following XML is an example of a consumer bucket XML file:

```

<?xml version="1.0" encoding="UTF-8"?>
<realtimedata messenger="ws" synchronous="true">

  <!-- ===== -->
  <!-- The attributes which are accepted by this consumer -->
  <!-- ===== -->

  <attributes>
    <attribute name="n1" type="number"/>
    <attribute name="n2" type="number"/>
    <attribute name="sum" type="number"/>
  </attributes>

  <!-- ===== -->
  <!-- The web service configuration -->
  <!-- ===== -->

  <messengerconfig>
    path = addup
  </messengerconfig>

  <!-- ===== -->
  <!-- The script which creates the output -->
  <!-- ===== -->

  <outgoing key="addup1">
    <messagebody>
      <script multirecord="true">
        <![CDATA[
          function build(recs, mtags, ctype) {
            var rec = recs[0]

            if (ctype == 'json') {
              return JSON.stringify({ inputs: [rec.n1, rec.n2], sum:
rec.sum }, null, 2)
            } else if (ctype == "dom") {
              return <response><inputs><v>{rec.n1}</v><v>{rec.n2}</v></
inputs><sum>{rec.sum}</sum></response>.toXMLString()
            } else {
              return rec.n1 + " + " + rec.n2 + " = " + rec.sum
            }
          }
        ]]>
      </script>
    </messagebody>
  </outgoing>

```

</realtimedata>