

Oracle® Fusion Middleware

Installing and Configuring Oracle Enterprise Data Quality

Release 12c (12.2.1)

E56640-01

October 2015

Documentation for installers and system administrators that describes how to install, configure, and upgrade Oracle Enterprise Data Quality.

Copyright © 2015 Oracle and/or its affiliates. All rights reserved.

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 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 installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. 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 and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon 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, Opteron, the AMD logo, and the AMD Opteron 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.

Contents

Preface	vii
Audience	vii
Documentation Accessibility	vii
Related Documents	vii
Conventions	viii
1 Planning an Enterprise Data Quality Installation	
1.1 Overview of EDQ	1-1
1.2 Overview of the Installation and Configuration Tasks	1-1
1.3 Choosing EDQ Components and Versions	1-2
1.3.1 Choosing the Correct Combination of EDQ Required Components	1-2
1.3.2 Supported Platforms and Component Versions	1-2
1.4 Satisfying EDQ System Requirements	1-3
1.4.1 Disk and Memory Requirements	1-3
1.4.2 EDQ Directory Requirements	1-3
1.4.3 UNIX System Resource Requirements (ulimit)	1-4
1.4.4 Virtual Hardware	1-4
1.4.5 Operating System User	1-5
1.5 Downloading EDQ	1-5
2 Installing the Required External Software Components	
2.1 Prerequisites for these Procedures	2-1
2.2 Installing a Java Development Kit to Support EDQ	2-1
2.3 Installing an Application Server to Support EDQ	2-2
2.3.1 Installing Oracle Fusion Middleware Infrastructure (includes Oracle WebLogic Server) 2-2	
2.3.2 Installing Apache Tomcat	2-3
2.4 Installing a Database to Support EDQ	2-3
2.4.1 Installing an Oracle Database to Support EDQ	2-4
2.4.1.1 Configuring Oracle Database to Support EDQ	2-4
2.4.2 Installing a PostgreSQL Database to Support EDQ	2-5
2.4.2.1 Configuring PostgreSQL Database to Support EDQ	2-5
2.4.2.2 Checking Local Database Connections	2-5
2.4.2.3 Enabling Password Authentication on Linux Systems	2-6
2.4.2.4 Configuring Remote Connections	2-6

3	Installing Enterprise Data Quality	
3.1	Starting the Installation Program	3-1
3.2	Navigating the EDQ Installation Screens	3-1
4	Configuring Enterprise Data Quality with Oracle WebLogic Server	
4.1	Prerequisites for these Procedures	4-1
4.2	Creating an EDQ Database Repository.....	4-1
4.3	Creating the WebLogic Server EDQ Domain	4-3
4.3.1	Starting the WebLogic Server Domain Configuration Wizard	4-4
4.3.2	Navigating the Domain Configuration Wizard Screens	4-4
4.3.3	Configuring Launchpad to Show the Managed Server	4-7
4.4	Start Oracle WebLogic Server	4-7
4.5	Running Multiple EDQ Servers in the Same Domain	4-7
4.6	Configuring EDQ for High-Availability in a WebLogic Server Cluster	4-8
4.6.1	Configuring EDQ for Clustered Mode Deployment	4-8
4.6.1.1	Running the Domain Configuration Wizard for a Clustered Deployment	4-8
4.6.2	Starting EDQ in a Cluster	4-9
4.6.3	Enabling JMX API and Command Line for HA Clusters	4-10
4.6.4	Landing Area.....	4-10
5	Configuring Enterprise Data Quality with Apache Tomcat	
5.1	Prerequisites for these Procedures	5-1
5.2	Creating the EDQ Database Objects.....	5-1
5.2.1	Creating an EDQ Repository in an Oracle Database	5-2
5.2.2	Creating an EDQ Repository in a PostgreSQL Database	5-2
5.3	Configuring Tomcat Application Server	5-2
5.4	Creating Directories for Use With Tomcat Application Server.....	5-3
5.5	Configuring EDQ to work with Tomcat Application Server	5-3
5.6	Verifying EDQ Functional Packs	5-4
5.7	Deploying the EDQ Application on a Tomcat Application Server	5-5
6	Setting Server Parameters to Support Enterprise Data Quality	
6.1	How to Set the Server Parameters	6-1
6.1.1	How to Set Server Parameters in a WebLogic Server Domain	6-1
6.1.2	How to Set Server Parameters in a Tomcat Domain	6-2
6.2	Recommended Parameter Values	6-2
6.2.1	Setting JVM Parameters	6-2
6.2.2	Setting Other Parameters.....	6-2
6.3	Example of Parameter Settings	6-2
7	Next Steps After Configuring Enterprise Data Quality	
7.1	Logging Into EDQ.....	7-1
7.2	The EDQ Launchpad	7-1

8 Upgrading Enterprise Data Quality

9 Removing Enterprise Data Quality from a System

9.1	Removing EDQ from a Linux or UNIX System	9-1
9.2	Removing EDQ from a Windows System	9-1

Preface

This document describes how to install and upgrade Oracle Enterprise Data Quality on Linux, UNIX, and Windows platforms. It also provides instructions for configuring EDQ to work with an application server and database.

Audience

This document is intended for system administrators or application developers who are installing Oracle Enterprise Data Quality. It is assumed that you have a basic understanding of application server and web technology and have a general understanding of Linux, UNIX, and Windows platforms. Throughout this guide, it is assumed that you are fully familiar with the components of the supported platform on which you want to install Oracle Enterprise Data Quality.

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>.

Access to Oracle Support

Oracle customers that 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 about EDQ, see the following documents in the Oracle Enterprise Data Quality documentation set.

EDQ Documentation Library

The following publications are provided to help you install and use EDQ:

- *Oracle Fusion Middleware Release Notes for Enterprise Data Quality*
- *Oracle Fusion Middleware Installing and Configuring Enterprise Data Quality*
- *Oracle Fusion Middleware Administering Enterprise Data Quality*
- *Oracle Fusion Middleware Understanding Enterprise Data Quality*
- *Oracle Fusion Middleware Integrating Enterprise Data Quality With External Systems*

- *Oracle Fusion Middleware Securing Oracle Enterprise Data Quality*
- *Oracle Enterprise Data Quality Address Verification Server Installation and Upgrade Guide*
- *Oracle Enterprise Data Quality Address Verification Server Release Notes*

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

<https://docs.oracle.com>

Online Help

Online help is provided for all Oracle Enterprise Data Quality 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.
monospace	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.

Planning an Enterprise Data Quality Installation

This chapter helps to prepare you for your Oracle Enterprise Data Quality (EDQ) installation. Various topics are covered that should be reviewed thoroughly to help ensure that you do not encounter any problems either during or after the product installation and domain configuration.

This chapter includes the following sections:

- [Overview of EDQ](#)
- [Overview of the Installation and Configuration Tasks](#)
- [Choosing EDQ Components and Versions](#)
- [Satisfying EDQ System Requirements](#)
- [Downloading EDQ](#)

1.1 Overview of EDQ

EDQ provides a comprehensive data quality management environment that is used to understand, improve, protect and govern data quality. The software facilitates best practice master data management, data integration, business intelligence, and data migration initiatives. It provides integrated data quality in customer relationship management (CRM) and other applications.

This documentation guides you through the selection, installation and configuration of the components that are needed to support EDQ.

1.2 Overview of the Installation and Configuration Tasks

This section ([Table 1-1](#)) provides an overview of the EDQ installation and configuration tasks that you will perform, in the order that they should be performed.

Table 1-1 EDQ Product Installation Procedure Tasks

Task	Action to Perform
Understand and select the external software components that support EDQ.	See Section 1.3, "Choosing EDQ Components and Versions."
Satisfy EDQ system requirements.	See Section 1.4, "Satisfying EDQ System Requirements"
Obtain an EDQ installation file from Oracle Software Delivery Cloud.	See Section 1.5, "Downloading EDQ."

Table 1–1 (Cont.) EDQ Product Installation Procedure Tasks

Task	Action to Perform
Install the JDK and your chosen application server and database components.	See Chapter 2, "Installing the Required External Software Components."
Install the EDQ software.	See Chapter 3, "Installing Enterprise Data Quality."
Configure EDQ	For instructions see, Chapter 4, "Configuring Enterprise Data Quality with Oracle WebLogic Server" or Chapter 5, "Configuring Enterprise Data Quality with Apache Tomcat."
Set system parameters.	See Chapter 6, "Setting Server Parameters to Support Enterprise Data Quality."
Next Steps	For log in and basic use information, see Section 7, "Next Steps After Configuring Enterprise Data Quality."
(Optional) Upgrade from a previous release of EDQ.	If using Oracle WebLogic Server or Apache Tomcat, see Section 8, "Upgrading Enterprise Data Quality."

1.3 Choosing EDQ Components and Versions

The following sections show you the components that are required to support EDQ and the supported versions of those components and EDQ.

- [Choosing the Correct Combination of EDQ Required Components](#)
- [Supported Platforms and Component Versions](#)

1.3.1 Choosing the Correct Combination of EDQ Required Components

EDQ is a Java Web Application that uses a Java Servlet Engine, a Java Web Start graphical user interface, and a data repository within a database. As such, it requires access to the following components:

- a Java Development Kit (JDK)
- a Java Application Server to supply web services. Oracle WebLogic Server and Apache Tomcat are supported.
- a structured query language (SQL) relational database management system (RDBMS) to store configuration data, working data, and the results of work performed by the processes. Oracle Database is supported.

The following application servers and databases are supported for use with EDQ:

- Oracle WebLogic 12.2.1 (from the Fusion Middleware Infrastructure 12.2.1 package), or Apache Tomcat 8.
- Oracle Database 11.2.0.4+ or 12.1.0.1+.

Instructions for installing these components are in [Section 2, "Installing the Required External Software Components."](#) See [Section 1.3.2](#) for supported versions of each of these components.

1.3.2 Supported Platforms and Component Versions

Review the list of certified platforms, JDKs, databases, application servers, and releases for EDQ prior to installation. This information is listed in *Oracle Enterprise Data Quality Certification Matrix* at

<http://www.oracle.com/technetwork/middleware/ias/downloads/fusion-certification-100350.html>

1.4 Satisfying EDQ System Requirements

This section describes the hardware and software requirements of EDQ. These requirements represent the server configurations that are certified and supported by Oracle for the EDQ product.

1.4.1 Disk and Memory Requirements

Depending on the tasks that EDQ is required to perform, it can place heavy demands on the hardware used to run it. A recommended *minimum* hardware specification for an EDQ server is:

- 8GB physical memory, with 4GB allocated to the EDQ Java Virtual Machine (JVM)
- At least 4 logical CPUs
- At least 500GB of hard disk space on the database server. The EDQ results schema, which contains the working data that EDQ generates, must have enough space to contain at least 20 times the volume of source data that you expect to process through EDQ. This size may increase if there are many EDQ users working on the same projects at the same time. The configuration schema remains small, normally less than 5GB, unless there is a large amount of user-modified reference data and Case Management data.

Note: These recommendations do not represent sizing advice for any specific deployment, but rather a starting point for testing size requirements in your environment. It may be appropriate to deploy a larger machine or many machines, depending on the processing load placed on EDQ.

1.4.2 EDQ Directory Requirements

EDQ uses an installation directory and two configuration directories. You should record the location of these directories in case you need to apply manual updates to any of their contents.

EDQ Installation Directory

During the installation process, you must specify an installation directory to contain the EDQ installation files. This directory is known as the EDQ Home (*EDQ_HOME*) directory and is named as follows:

- If you are installing EDQ as part of the Oracle Fusion Middleware product stack, you must install the EDQ Home directory as a subdirectory of the Oracle Fusion Middleware home (installation) directory. A typical default Fusion Middleware home directory is as follows, depending on the platform:

Linux and UNIX:

```
/opt/Oracle/Middleware/FMW_HOME
```

Windows:

```
C:\Oracle\Middleware\FMW_HOME
```

Note: The Middleware home directory is referenced as *FMW_HOME* in this guide.

- If you use the WebLogic Server or Apache Tomcat application server, you can install the EDQ Home directory in any directory.

The EDQ Home directory requires approximately 1GB of hard disk space.

EDQ Configuration Directories

EDQ requires two configuration directories, which are separate from the EDQ Home (installation) directory that contains the program files. The configuration directories are:

- The *base* configuration directory: This directory contains default configuration data. Once EDQ is installed, the files in the base configuration directory must not be altered, renamed, or moved.
- The *local* configuration directory: This directory contains overrides to the default configuration. EDQ looks for overrides in this directory first, before looking in the base configuration directory. Files in the local configuration directory can be modified to customize or extend EDQ.

The names and locations of the configuration directories are as follows:

- If you are using Oracle WebLogic Server, the Configuration Wizard automatically creates and populates the configuration directories in the EDQ domain with the names of `oedq.home` (base configuration directory) and `oedq.local.home` (local configuration directory). An example installation path is:

```
WLS_HOME/user_projects/domains/edq/config/fmwconfig/edq/oedq.home  
WLS_HOME/user_projects/domains/edq/config/fmwconfig/edq/oedq.local.home
```

- If you are using Apache Tomcat, you create the configuration directories manually in any location, with any names, and the configuration utility will populate them. You are prompted to create the directories during the installation instructions.

On a default EDQ installation, the configuration directories occupy approximately 1MB in total. When an EDQ instance uses the `landingarea`, this total can increase as files are loaded up for processing and written out.

1.4.3 UNIX System Resource Requirements (ulimit)

Depending on how the UNIX system that hosts the application server is configured, you may find that the application server cannot create files larger than 1 GB. This restricts your ability to work with large data sets if you intend to use files to transfer data to EDQ for processing. The most common limit that has to be changed is the number of processes that the application server user (for example, 'weblogic') can create. On large servers, this often needs to be increased, for example to 65536.

System resource limits are controlled by the `ulimit` command. Default `ulimit` values exist, which you can view by using the `ulimit -a` command. Look at the settings for file size, process limit, and file handle limits. The hard `ulimit` on file size may need to be adjusted upward or removed for your application server account. Consult your System Administrator for assistance, if needed.

1.4.4 Virtual Hardware

You can install EDQ on virtualized systems using a virtualization tool, such as Oracle VM Server. Both the virtual system and the physical system must fulfill the minimum hardware requirements listed in this documentation.

If load balancing software is used to deploy multiple virtual systems onto a single physical system, care must be taken to ensure that the load balancing software is carefully tuned. In general, EDQ requires considerable compute power and memory in the Middle Tier (where it performs most of its processing), and considerable tablespace in the Database (where it writes data and results). Between batches, very little load is imposed on the system. When processing a batch of data, EDQ rapidly drives hardware to be CPU or I/O bound. Unless the virtualized load balancing is correctly configured, suboptimal performance results.

1.4.5 Operating System User

An operating system user account is used to install and upgrade EDQ on your servers and install the application server. This user is required on all platforms. This user *must* have full permissions (read, write and execute) to the EDQ installation (*EDQ_HOME*) directory, configuration directories, and all database directories. This user account is referred to as the *EDQ installation user* in this documentation. For more information about the EDQ directories, see [Section 1.4.2, "EDQ Directory Requirements."](#)

Note: When installing on UNIX or Linux operating systems, do *not* use the root user as your EDQ installation user account.

1.5 Downloading EDQ

To download the EDQ installation and configuration files, obtain the generic package installer from the Oracle Software Delivery Cloud website as follows:

1. Enter the Oracle Software Delivery Cloud URL into a web browser:

<http://edelivery.oracle.com/>

2. Click **Sign-in/Register**.

Note: If you are not already logged in, the Oracle Single Sign-On page appears. Enter your Oracle user id and password and click **Sign In**.

The Terms & Restrictions page appears.

3. Select the **Oracle Software Delivery Cloud Trial License Agreement** and the **Export Restrictions** check boxes, and then click **Continue**.

The Media Pack Search page appears.

4. On the Media Pack Search page, do the following:
 - a. From the **Select Product Pack** drop-down list, select **Oracle Fusion Middleware**.
 - b. From the **Platform** drop-down list, select the platform on which you are installing EDQ.
 - c. Click **Go**.

The Results list expands to show all available media packs that include your search criteria.

5. Locate and select **Oracle Enterprise Data Quality 12c (12.2.1) Media Pack** (Oracle Fusion Middleware Product Pack), and then click **Continue**.

6. Click the **Download** button next to **Oracle Enterprise Data Quality 12.2.1**.
7. Browse to the directory where you want to save the file. Click **Save** to start the file download. A ZIP file is downloaded.
8. Extract the ZIP file to a temporary directory.

Installing the Required External Software Components

This chapter describes how to install the software components that support EDQ.

This chapter includes the following sections:

- [Prerequisites for these Procedures](#)
- [Installing a Java Development Kit to Support EDQ](#)
- [Installing an Application Server to Support EDQ](#)
- [Installing a Database to Support EDQ](#)

Note: These components must be installed prior to installing EDQ.

2.1 Prerequisites for these Procedures

Before performing the procedures in this section, you must first read and satisfy the requirements in [Chapter 1, "Planning an Enterprise Data Quality Installation."](#)

2.2 Installing a Java Development Kit to Support EDQ

EDQ and the application server both rely on the Java Development Kit (JDK). The JDK provides a Java run-time environment (JRE) and tools for compiling and debugging Java applications.

Use the following table to identify the appropriate JDK to install on the platform where you are installing EDQ.

Platform	JDK to Use
AIX	IBM JDK. This is the only supported JDK on AIX.
HP-UX	HP JDK
Solaris	Oracle JDK
Linux	Oracle JDK
Windows	Oracle JDK

For supported JDK versions per platform, see the *Oracle Enterprise Data Quality Certification Matrix* at

<http://www.oracle.com/technetwork/middleware/ias/downloads/fusion-certification-100350.html>

Download and install the appropriate JDK using the instructions provided at <http://www.oracle.com/technetwork/java/javase/downloads/index.html>

Note: On Solaris systems, you must install both the 32-bit and 64-bit JDKs in order to run Java applications.

Make a note of the directory into which you installed the JDK. You will need to specify this directory during the installation of the application server. The path to this directory is referred to as the *JDK_HOME* directory in this documentation. For example, the *JDK_HOME* may be like one of the following:

Linux and UNIX:

/opt/jdk1.8

Windows:

C:\Program Files\Java\jdk1.8

Note: If your *JDK_HOME* directory contains spaces, as in the default installation directory of C:\Program Files\Java\jdk1.8, you must set your *JAVA_HOME* environment variable to the Windows short name for the directory, as follows: C:\Progra~1\Java\jdk1.8.

2.3 Installing an Application Server to Support EDQ

For maximum security, you should configure the application server (whether Oracle WebLogic Server or Apache Tomcat) to use HTTPS only. You should disable unencrypted HTTP access. You can alternatively use a frontend such as Oracle HTTP Server or Apache HTTP Server, where HTTPS is provided.

Choose and download one of the following application servers:

- [Installing Oracle Fusion Middleware Infrastructure \(includes Oracle WebLogic Server\)](#)
- [Installing Apache Tomcat](#)

Note: The application server that you choose determines the database that you can use. With Oracle WebLogic Server, you must use Oracle Database, and you must install EDQ within an Oracle Fusion Middleware environment. With Apache Tomcat, you can use either Oracle Database or PostgreSQL Database, and you can install EDQ as a standalone application. PostgreSQL is not supported for production use.

2.3.1 Installing Oracle Fusion Middleware Infrastructure (includes Oracle WebLogic Server)

When installing on Oracle WebLogic Server, install both Oracle WebLogic Server and JRF. To view the WebLogic Server versions that are supported by EDQ in this release, see Oracle Enterprise Data Quality Certification Matrix at

<http://www.oracle.com/technetwork/middleware/ias/downloads/fusion-certification-100350.html>

To download and install Oracle WebLogic Server and JRF, see *Installing and Configuring Oracle WebLogic Server and Coherence*.

Note: Use the Fusion Middleware Infrastructure installer. Do not use the free WebLogic Server distribution for developers, because this installer does not contain all of the required files for EDQ.

The directory path to your WebLogic Server installation is referenced as the `WLS_HOME` directory in this documentation.

2.3.2 Installing Apache Tomcat

To view the Tomcat versions that are supported by EDQ in this release, see the Oracle Enterprise Data Quality Certification Matrix at

<http://www.oracle.com/technetwork/middleware/ias/downloads/fusion-certification-100350.html>

To download and install Apache Tomcat, go to the Apache Software Foundation Server website at

<http://tomcat.apache.org>

After you install Tomcat, perform the following tasks:

1. Configure Tomcat to use the Java Development Kit (JDK) that you installed in [Section 2.2, "Installing a Java Development Kit to Support EDQ."](#) You can do one of the following:
 - Set the path to the JDK executable in the `/etc/sysconfig/tomcat7.conf` file. For example:


```
JAVA_HOME="/opt/java/jdk1.7.0_45"
```
 - Add the JDK path to the local `setenv.sh` file.
2. (Recommended) Configure Tomcat to start as a service.
3. Create an operating system user who will run Tomcat. This user must own the EDQ installation and configuration directories. This account is used to administer your EDQ domain and to log into the EDQ application. For more information about the EDQ directories, see [Section 1.4.2, "EDQ Directory Requirements."](#)

2.4 Installing a Database to Support EDQ

This section contains information that is specific to the installation and configuration of the database that will contain the EDQ repository.

- If using Oracle WebLogic Server as the EDQ application server, you must install Oracle Database.
- If using Apache Tomcat as the application server, the supported database for production is Oracle Database. PostgreSQL database is supported only for development or test purposes.

To view the supported database versions, see

<http://www.oracle.com/technetwork/middleware/ias/downloads/fusion-certification-100350.html>

Note: You will create three schemas and three users in this database when configuring EDQ in [Section 4, "Configuring Enterprise Data Quality with Oracle WebLogic Server"](#) or [Section 5, "Configuring Enterprise Data Quality with Apache Tomcat"](#) (depending on the application server you chose).

2.4.1 Installing an Oracle Database to Support EDQ

You can download a supported Oracle Database product and installation instructions from the Oracle Database website at

<http://www.oracle.com>

2.4.1.1 Configuring Oracle Database to Support EDQ

The following configuration elements are either required or recommended when installing or configuring an Oracle database for use with EDQ.

Required:

- Select the **Create and configure a database** installation option.
- Configure the database to use a Unicode character set to ensure that EDQ is able to capture and process data in the widest range of character sets. For more information, see "Supporting Multilingual Databases with Unicode" in *Oracle Database Globalization Support Guide*.
- EDQ requires a database administrator (DBA) user account in the database. This database account is used to access the database during the installation and configuration processes to create database accounts and objects that are specific to EDQ.

Recommended:

- Oracle recommends making the following configuration selections when specifying the Oracle memory structure and tablespace configuration to support EDQ:
 - 4GB Program Global Area (PGA)
 - 4GB System Global Area (SGA)
 - 20GB undo tablespace
 - 20GB temp tablespace
 - 120GB tablespace (4 x 30GB files) for Results database (EDQRESULTS schema)
 - 30GB (1 x 30GB file) for Config database (EDQCONFIG schema)
- You may need to increase the values for the `SESSIONS` and `PROCESSES` parameters. The suggested values are as follows, but these parameters may need to be adjusted later for optimal performance:

```
SESSIONS=500
```

```
PROCESSES=500
```

For more information about setting these values appropriately, contact your database administrator.

2.4.2 Installing a PostgreSQL Database to Support EDQ

You can download the PostgreSQL product and installation instructions from the official PostgreSQL website at

<http://www.postgresql.org>

Note: If PostgreSQL was distributed with your operating system, verify that the release is one of EDQ supported releases by going to the certification matrix at

<http://www.oracle.com/technetwork/middleware/ias/downloads/fusion-certification-100350.html>

2.4.2.1 Configuring PostgreSQL Database to Support EDQ

The following configuration elements are either required or recommended when installing or configuring a PostgreSQL database for use with EDQ.

Required:

- Allow 403 maximum connections by editing the `postgresql.conf` file in the PostgreSQL data directory (for example, `var/lib/pgsql/data/postgresql.conf`).
- Configure the PostgreSQL database to use a Unicode character set to ensure that EDQ is able to capture and process data in the widest range of character sets. For more information, click the **Documentation** link on the PostgreSQL website at <http://www.postgresql.org>
- To use a PostgreSQL Database with EDQ, you must enable database listening. See [Section 2.4.2.2, "Checking Local Database Connections."](#)
- On Linux systems, you must configure PostgreSQL to use password authentication. See [Section 2.4.2.3, "Enabling Password Authentication on Linux Systems."](#)

Recommended:

If you are installing on Windows, Oracle recommends that you use the graphical installer that you can download from the PostgreSQL website at

<http://www.postgresql.org>

Each operating system on the Download page provides a link to the graphical installer.

2.4.2.2 Checking Local Database Connections

This section describes how to ensure that the PostgreSQL database local connections are properly configured to ensure database listening is on and that passwords are accepted for authentication.

1. Log in to the system as the EDQ database user.
2. (Windows only) Locate the MS-DOS Command Prompt (`cmd.exe`), and then run it.
3. Ensure that your database server is running. For more information, see the PostgreSQL documentation for your installed version on the PostgreSQL website at

<http://www.postgresql.org>

4. Go to the PostgreSQL data directory. For example, `/var/lib/pgsql/data` on Linux or `C:\Program Files\PostgreSQL\9.3\data` on Windows.
5. Open the `pg_hba.conf` file by entering the one of the following commands, depending on the platform:

On the Linux command line:

```
tail pg_hba.conf
```

On the Windows command line:

```
TYPE pg_hba.conf
```

The file contents are displayed and the following is an excerpt:

```
# TYPE DATABASE USER ADDRESS METHOD
# "local" is for Unix domain socket connections only
local all all ident

# IPv4 local connections:
host all all 127.0.0.1/32 md5
# IPv6 local connections:
host all all ::1/128 md5
```

6. Ensure that the IPv4 and IPv6 local connections are configured as shown in the console output in the previous step. Use the address for your own database in the ADDRESS column.

2.4.2.3 Enabling Password Authentication on Linux Systems

On Linux systems, you must configure PostgreSQL to use password authentication by editing the `pg_hba.conf` file in the data subdirectory of the PostgreSQL installation directory (for example, `/var/lib/pgsql/data/pg_hba.conf`). Locate the `ident sameuser` entries, and then change them to a value of `md5`.

2.4.2.4 Configuring Remote Connections

If your PostgreSQL database is installed on a different system from the one where EDQ is installed, it must be configured to accept connections from other hosts. This is applicable to all PostgreSQL environments.

1. Locate the data directory for your PostgreSQL database.
2. Edit the `postgresql.conf` file.
3. Locate the following line:

```
#listen_addresses = 'localhost'
```

4. Insert the following line to cause PostgreSQL to accept connections from remote hosts:

```
listen_addresses = '*'
```

5. Edit the `pg_hba.conf` file then add the following line using `addr/mask` to identify the subnet of the host running EDQ:

```
host all all addr/mask md5
```

For example, `host all all 192.168.0.0/24 md5` allows connections from all hosts with the IP addresses 192.168.0.0 to 192.168.0.255 while `host all all 0.0.23.56/32 md5` accepts connections only from IP address 10.0.23.56.

Installing Enterprise Data Quality

This chapter describes how to start the Enterprise Data Quality installation program in graphical mode. It also describes the sequence of screens that appear in the installation process.

This chapter includes the following sections:

- [Starting the Installation Program](#)
- [Navigating the EDQ Installation Screens](#)

3.1 Starting the Installation Program

The EDQ installation program is delivered as a generic Java Archive (JAR) file that is used to install all supported operating systems.

To start the installation program, perform the following steps.

1. Log in to the target system as your EDQ installation user. This user must be able to run the command prompt as an administrator.
2. Go to the directory where you downloaded and unzipped the installation program.
3. Launch the installation program by invoking `java -jar` from the JDK installation directory (*JDK_HOME*) directory on your system as follows:

On UNIX operating systems:

```
JDK_HOME/bin/java -jar edq_generic.jar
```

On Windows operating systems:

```
JDK_HOME\bin\java -jar edq_generic.jar
```

In the above commands for UNIX and Windows operating systems, `edq_generic.jar` is a generic place-holder for the actual name of the installation jar file. When the installation program appears, you are ready to begin the installation. See [Section 3.2](#) for a description of each installation program screen.

3.2 Navigating the EDQ Installation Screens

The installation program displays a series of screens. If you need additional help with an installation screen, click **Help** on the installation screens during installation. Complete the installation using the instructions in [Table 3-1, "Running the Installation Program"](#)

Table 3–1 Running the Installation Program

Screen	Action to Perform
Installation Inventory Setup	<p>On UNIX operating systems, this screen opens if this is the first time you are installing an Oracle product on this host. Specify the location where you want to create your central inventory. Make sure that the operating system group name selected on this screen has write permissions to the central inventory location.</p> <p>For more about the central inventory, see "Understanding the Oracle Central Inventory" in <i>Installing Software with the Oracle Universal Installer</i>.</p> <p>This screen does not appear on Windows operating systems.</p>
Welcome	<p>Click Next to proceed with the installation.</p> <p>You may cancel the installation at any time by clicking Exit.</p>
Installation Location	<p>Specify the Middleware installation directory (<code>FMW_HOME</code>), which will contain EDQ once installed. See Section 1.4.2, "EDQ Directory Requirements."</p> <p>Click Next to continue.</p>
Prerequisite Checks	<p>This screen verifies that your system meets the minimum necessary requirements.</p> <p>Click Next to continue.</p>
Installation Summary	<p>Use this screen to verify the installation options you selected. If you want to save these options to a response file, click Save Response File and provide the location and name of the response file. Response files can be used later in a silent installation situation.</p> <p>Click Install to continue.</p>
Installation Progress	<p>This screen allows you to see the progress of the installation.</p> <p>When the installation program progress reaches 100%, click Finish to dismiss the installer or click Next to see a summary.</p>
Installation Complete	<p>This screen appears when the installation is complete.</p> <p>Click Finish to exit the installation program.</p>

Configuring Enterprise Data Quality with Oracle WebLogic Server

This chapter describes how to create the EDQ database repository, create an Oracle WebLogic Server domain for EDQ, and start Oracle WebLogic Server.

Note: These instructions apply to Oracle WebLogic Server environments only. If you are using Apache Tomcat, you must follow the directions in [Chapter 5, "Configuring Enterprise Data Quality with Apache Tomcat."](#)

This chapter includes the following sections:

- [Prerequisites for these Procedures](#)
- [Creating an EDQ Database Repository](#)
- [Creating the WebLogic Server EDQ Domain](#)
- [Start Oracle WebLogic Server](#)
- [Running Multiple EDQ Servers in the Same Domain](#)
- [Configuring EDQ for High-Availability in a WebLogic Server Cluster](#)

4.1 Prerequisites for these Procedures

Before performing the procedures in this section, you must first read and satisfy the steps in:

- [Chapter 1, "Planning an Enterprise Data Quality Installation"](#)
- [Chapter 2, "Installing the Required External Software Components"](#)
- [Chapter 3, "Installing Enterprise Data Quality"](#)

4.2 Creating an EDQ Database Repository

EDQ makes use of some database schemas. These schemas are the configuration schema (EDQCONFIG), the results schema (EDQRESULTS), and the staging schema (EDQSTAGING). You create them with the Oracle Repository Creation Utility (RCU).

The person who runs RCU must be able to log into the database with DBA privileges. If you cannot run with DBA privileges, RCU can create a script for a DBA to run later.

Note: Do not use RCU to upgrade EDQ; use the instructions in [Section 8, "Upgrading Enterprise Data Quality."](#)

To run RCU:

1. Make certain the repository database is running.
2. Run the command shell or console of the operating system.
3. Start RCU from the `FMW_HOME/oracle_common/bin` directory, where `FMW_HOME` is the Oracle Fusion Middleware installation directory.

On Linux:

```
./rcu
```

On Windows:

```
rcu.bat
```

Complete the RCU configuration screens by following the instructions in [Table 4-1](#).

Table 4-1 *Running the RCU Program*

Screen	Action to Perform
Welcome	Click Next to proceed with the installation. You may cancel the installation at any time by clicking Exit .
Create Repository	Click Next to continue. This uses the default Create Repository and System Load and Product Load options. This requires the person running RCU to have DBA privileges.
Database Connection Details	Select Oracle Database from the Database Type list. Specify the host name where your Oracle database is running. Enter the port number for your database. The default port number for Oracle Database is 1521. Specify the service name for the database. Typically, the service name is the same as the global database name. For example, <code>orcl.example.com</code> . Note that if you are using release 12c of the Oracle Database, you need to ensure that you enter the connection details of a pluggable database. For example, <code>pdborcl.example.com</code> . Enter the user name for your database. The user name could either be <code>SYS</code> or that of the user with DBA privileges. Enter the password for your database user. Select SYSDBA from the Role: list. This is automatically selected when the user is <code>SYS</code> . Select Normal from the Role: list if you are a user with DBA privileges. Click Next to continue.
Checking Global Prerequisites	When the prerequisites checking progress has reached completion, click OK to continue.

Table 4–1 (Cont.) Running the RCU Program

Screen	Action to Perform
Select Components	<p>Select Create new prefix and enter a unique prefix name for all the database schemas you are creating in this session. For example, <code>edqprod</code> or the default of <code>DEV</code>.</p> <p>Select the Oracle AS Repository Components check box. The Oracle EDQ check boxes that create the EDQ configuration, results, and staging schemas in the database repository are then automatically selected.</p> <p>When you select the Oracle EDQ component, all the necessary schemas are installed. Not only will the EDQ schema be selected but the common schema needed to support audit and OPSS, which EDQ is integrated with will also be selected. This ensure that EDQ will function correctly with these integrated components.</p> <p>The three EDQ schemas that are to be installed are Config, Results and Staging. The Staging schema is also used by CDS, with or without the Fusion Connector. However installing the Staging schema when not necessary will not interfere with operation and will consume minimal database resources.</p> <p>Click Next to continue.</p>
Checking Component Prerequisites	<p>When the prerequisites checking progress has reached completion, click OK to continue.</p>
Schema Passwords	<p>Ensure that Use same passwords for all schemas is selected.</p> <p>Enter the password that you want to use for all of the EDQ database schemas in all password fields, then click Next to continue.</p>
Map Tablespaces	<p>The default EDQ tablespaces that will be created by RCU are displayed by component.</p> <p>You can change the tablespaces by clicking Manage Tablespaces and then modifying the information. Oracle recommends using one user tablespace for the <code>EDQCONFIG</code> schema, and a different user tablespace for the <code>EDQRESULTS</code> schema. The recommended minimum sizes (can be adjusted later) are:</p> <ul style="list-style-type: none"> ■ <code>EDQRESULTS</code>: 120GB tablespace (4 x 30GB files) ■ <code>EDQCONFIG</code>: 30GB tablespace (1 x 30GB file) ■ <code>EDQSTAGING</code>: 30GB tablespace (1 x 30GB file) <p>Click Next to continue.</p>
Validating and Creating Tablespaces	<p>Click OK to create any non-existent tablespaces in your schema, then click OK when the operation completes.</p>
Summary	<p>Review the database details, then click Create to continue.</p> <p>A status screen is displayed that shows the progress of creating the repository components.</p>
Completion Summary	<p>Click Close to exit the RCU program.</p>

4.3 Creating the WebLogic Server EDQ Domain

These instructions use the Configuration Wizard to create a Basic WebLogic Server domain for EDQ, with the following:

- One Administration Server and one managed server (no additional managed servers or clusters).

Note: If there is the possibility you may add more managed servers in the future, you should choose a cluster deployment, even if it starts with a single server. Otherwise it will be necessary to manually re-target the Data Sources, Applications and Deployments in WebLogic from a single server to the cluster.

- One (non-RAC) data source for the EDQ configuration schema, one data source for the results schema, and one data source for the staging schema. You can convert the datasources to RAC data sources with the Configuration Wizard, or you can do so later through the WebLogic Server Administration Console.
- A Node Manager configuration that is predefined within the EDQ domain as `edq/nodemanager`. You cannot edit the Node Manager home in this configuration. You can change this configuration during this procedure, if desired.

Note: Oracle recommends the use of managed servers that are administered by Oracle WebLogic Node Manager. You can configure Managed Servers, Clusters, and other advanced features through the Configuration Wizard, but it may be more practical to do so by using the WebLogic Server Administration Console after the initial configuration process. For more information, see [Section 4.5, "Running Multiple EDQ Servers in the Same Domain."](#)

4.3.1 Starting the WebLogic Server Domain Configuration Wizard

To start the Domain Configuration wizard, follow these steps. You will run the configuration wizard in graphical mode.

1. Log in to the system as the EDQ installation user that you created in [Section 1.4.5, "Operating System User."](#)
2. Go to `FMW_HOME/oracle_common/common/bin` directory, where `FMW_HOME` is the Fusion Middleware installation directory.
3. Start the wizard by entering the following command:

On Linux or UNIX operating systems:

```
./config.sh
```

On Microsoft Windows operating systems:

```
config.cmd
```

The WebLogic Server Configuration Wizard is displayed.

4.3.2 Navigating the Domain Configuration Wizard Screens

[Table 4–2](#) describes the screens in the configuration wizard. Certain screens are displayed only in certain situations depending on your selections. For help with any screen, click the **Help** button.

Table 4–2 Configuration Screens for Creating a New EDQ WebLogic Server Domain

Screen	Action to Perform
Create Domain	<p>Select Create a new domain.</p> <p>In the Domain Location box, enter the path to the new domain (for example, <i>FMW_HOME/user_projects/domains/edq_domain</i>) or click Browse to create the domain directory.</p> <p>Click Next to continue.</p>
Templates	<p>Select Oracle Enterprise Data Quality – 12.2.1.0 [edq]. The Oracle JRF and WebLogic Coherence Cluster Extension are automatically selected. Keep these selections.</p> <p>Click Next to continue.</p>
Administrator Account	<p>Specify the user name and password for the EDQ domain's administrator account. This account is used to administer the domain and to log into the EDQ application.</p> <p>Click Next to continue.</p>
Application Location	<p>Specify the directory in which the applications of the EDQ domain are to be stored.</p> <p>Click Next to continue.</p>
Domain Mode and JDK	<p>Domain Mode: Select the startup operation mode for your domain from the following options:</p> <ul style="list-style-type: none"> ▪ Development Mode—In this mode, the security configuration is relatively relaxed and you utilize <code>boot.properties</code> for the username and password, allowing you to auto-deploy applications. ▪ Production Mode—In this mode, the security configuration is relatively stringent, requiring a username and password to deploy applications. Before putting a domain into production, familiarize yourself with procedures for securing a production environment. For more information, see <i>Securing a Production Environment for Oracle WebLogic Server</i>. <p>JDK: From the Available JDKs list, select the JDK that you installed in Section 2.2, "Installing a Java Development Kit to Support EDQ."</p> <p>Click Next to continue.</p>
Database Configuration Type	<p>Ensure that RCU Data is selected. This populates the connection information you supplied when you ran the Repository Creation Utility (see Section 4–1, "Running the RCU Program").</p> <p>If you must change any of these fields, ensure that you use the schema prefix (<code>DEV</code> by default) and password that you specified when you ran RCU.</p> <p>When done, click Get RCU Configuration to connect to the Oracle Database and bind the EDQ schemas.</p> <p>Click Next to continue.</p>
Component Datasources	<p>Accept the defaults and then click Next.</p>
JDBC Test	<p>All schemas are selected and automatically tested.</p> <p>Return to the previous screen to alter the connection configuration if necessary.</p> <p>Click Next to continue.</p>
Advanced Configuration	<p>Select the Administration Server, Node Manager, Managed Servers, Clusters and Coherence Clusters option.</p>

Table 4–2 (Cont.) Configuration Screens for Creating a New EDQ WebLogic Server

Screen	Action to Perform
Administration Server	<p>On the Administration Server screen, change the listen address from "All Local Addresses" to the IP address of the host where the Administration Server will reside.</p> <p>Do not use All Local Addresses.</p> <p>Do not specify any server groups for the Administration Server.</p>
Node Manager	<p>Let the Per Domain Default Location option remain selected, and enter a Username and Password for the Node Manager.</p>
Managed Servers	<p>Clone the EDQ server to create a copy of the server.</p> <p>If you do not want to clone the EDQ server, select Add to add additional EDQ servers.</p> <p>Select the IP address of the host on which the Managed Server will reside.</p> <p>For configuration procedures for a clustered mode installation, see Section 4.6, "Configuring EDQ for High-Availability in a WebLogic Server Cluster"</p>
Clusters	<p>Select Add to add a cluster. If both servers are being deployed on the same machine, enter the machine name as the cluster address.</p> <p>For additional configuration procedures for a clustered mode installation and a description of those configuration screens, see Section 4.6, "Configuring EDQ for High-Availability in a WebLogic Server Cluster"</p> <p>To deploy EDQ in a non-clustered mode, select the appropriate components as needed.</p>
Assign Servers to Clusters	<p>Assign the managed servers that you created to the cluster that you created.</p> <p>For configuration procedures for a clustered mode installation, see Section 4.6, "Configuring EDQ for High-Availability in a WebLogic Server Cluster"</p>
Coherence Clusters	<p>Configure the Coherence cluster that is automatically added to the domain.</p> <p>Leave the default port number 0 as the Coherence cluster listen port.</p>
Machines	<p>Create a new machine in the domain. A machine is required so that the Node Manager can start and stop servers.</p> <p>Click Add to create a new machine.</p> <p>Specify <code>edq_machine_1</code> in the Name field.</p> <p>In the Node Manager Listen Address field, select the IP address of the machine in which the Managed Servers are being configured. You must select a specific interface and not "localhost." This allows Coherence cluster addresses to be dynamically calculated.</p> <p>Verify the port in the Node Manager Listen Port field. The port number 5556, shown in this example, may be referenced by other examples in the documentation. Replace this port number with your own port number as needed.</p>

Table 4–2 (Cont.) Configuration Screens for Creating a New EDQ WebLogic Server

Screen	Action to Perform
Assign Servers to Machines	<p>Assign the Administration Server and Managed Servers to the new machine you just created.</p> <p>In the Machines pane, select the machine you want to assign the servers to; in this case, <code>edq_machine_1</code>.</p> <p>In the Servers pane, assign <code>AdminServer</code> to <code>edq_machine_1</code> by doing one of the following:</p> <ul style="list-style-type: none"> ▪ Click once on <code>AdminServer</code> to select it, then click on the right arrow to move it beneath the selected machine (<code>edq_machine_1</code>) in the Machines pane. ▪ Double-click on <code>AdminServer</code> to move it beneath the selected machine (<code>edq_machine_1</code>) in the Machines pane. <p>Repeat to assign both <code>edq_server_1</code> and <code>edq_server_2</code> to <code>edq_machine_1</code>.</p>
Configuration Summary	<p>Review the configuration for your domain by selecting a view and then selecting individual items in the list for that view.</p> <p>If the domain is configured as you want it, click Create to create the domain.</p> <p>If you need to make changes to the configuration, click Back to return to the appropriate screen for the settings you want to change, or click on the links on the left to go to that screen.</p>
Configuration Progress	<p>Shows the progress of the domain creation.</p> <p>When the process completes, click Next.</p>
Configuration Success	<p>Review the domain creation results.</p> <p>Click Finish to exit the Configuration Wizard.</p>

4.3.3 Configuring Launchpad to Show the Managed Server

To configure a launchpad to show the managed server that it is connected to, add this line to the `director.properties` in the local home directory:

```
[expr]adf.headerextra = ': ' || weblogic.Name
```

4.4 Start Oracle WebLogic Server

You must start your Administration Server, Managed Servers, and clusters to complete the installation. For information about starting managed servers using Node Manager and Administration Servers, see "Starting and Stopping Oracle WebLogic Server Instances" in *Oracle Fusion Middleware Administering Server Startup and Shutdown for Oracle WebLogic Server*.

See also [Section 6, "Setting Server Parameters to Support Enterprise Data Quality"](#) for important information about setting server parameters for startup.

4.5 Running Multiple EDQ Servers in the Same Domain

To support high availability scenarios, Oracle recommends that you configure a cluster of multiple EDQ servers to share the incoming load (for example, from a large number of simultaneous web service requests), and to provide continuous service in the event of failure of an individual server. This section provides some basic guidance about how to configure EDQ to support such a model using Oracle WebLogic Server.

Multiple EDQ managed servers can be configured to run in the same WebLogic Server domain either in a cluster or not. If all the servers are on the same machine, each server must listen on a different port.

The Java Required Files (JRF) Template *must* be applied to any managed servers that were created using the WebLogic Server Administration Console. This is equivalent to the library targeting performed automatically by the WebLogic Server Configuration Wizard.

The final step is:

- Use the WebLogic Server Administration Console to modify the managed server settings for the additional EDQ servers.

Once multiple EDQ servers have been configured, you can leave them un-clustered and accessed directly using their respective Launchpad URLs to the relevant port, or you can configure them as part of a cluster using standard WebLogic Server practices. You can configure a separate front-end load balancer to handle incoming web service requests through a single cluster URL.

4.6 Configuring EDQ for High-Availability in a WebLogic Server Cluster

You can install and configure Oracle Enterprise Data Quality for high availability in an Oracle WebLogic Server cluster environment. The high availability features in Oracle Enterprise Data Quality have been enhanced to make the system function in a clustered environment, and tolerate individual RAC node failures and reconnect after complete database failures.

For more information on the high-availability features in Enterprise Data Quality, please see *Oracle Fusion Middleware Understanding Oracle Enterprise Data Quality*.

4.6.1 Configuring EDQ for Clustered Mode Deployment

To install and configure Oracle Enterprise Data Quality for high availability in an Oracle WebLogic Server cluster environment:

- Plan the installation, and install the prerequisites and software, as described in [Chapter 4.1, "Prerequisites for these Procedures"](#).
- Run the Repository Creation Utility (RCU), as described in [Section 4.2, "Creating an EDQ Database Repository"](#). To get a full high availability implementation, you must deploy the schema to an Oracle RAC instance. Deploying to a standard Oracle database does not provide HA capability for the database.
- Run the Domain Configuration wizard to provision a domain, as described in [Section 4.3, "Creating the WebLogic Server EDQ Domain"](#) and with the configuration steps for a cluster noted in [Section 4.6.1.1, "Running the Domain Configuration Wizard for a Clustered Deployment"](#).
- Start EDQ in a cluster as described in [Section 4.6.2, "Starting EDQ in a Cluster"](#).
- Enable or configure options for high availability in EDQ. See the sections:
 - [Section 4.6.3, "Enabling JMX API and Command Line for HA Clusters"](#)
 - [Section 4.6.4, "Landing Area"](#)

4.6.1.1 Running the Domain Configuration Wizard for a Clustered Deployment

To deploy EDQ in a clustered mode using the schema you created in the RCU:

1. Follow the steps in [Section 4.3.1, "Starting the WebLogic Server Domain Configuration Wizard"](#) and [Section 4.3.2, "Navigating the Domain Configuration Wizard Screens"](#) to launch and configure the basic configuration screens in the Domain Configuration wizard.
2. On the Advanced Configuration screen, select the **Managed Servers, Clusters and Coherence** option and then select **Next**.
3. On the Manage Servers screen, select the EDQ server, for example "edq_server1", and select the **Clone** button to create a copy of the server. If you do not want to clone the server, select the **Add** button to add a new server and target it to the cluster.

Name this server appropriately, for example "edq_server2".

The listen port should not be the same port as in Server 1 and the Server Group should be EDQ Managed Server. The listen address should be the same as server 1, assuming they are both running on the same machine.

Click **Next** and move to the Clusters screen.

4. Select the **Add** button to add a cluster.

If servers are being deployed on the same machine, enter the machine name as the cluster address. If not, the cluster address should be a comma-separated list of the machine names.

Leave the frontend host blank if you are not deploying a load balancer. If you are deploying with a load balancer, enter the URL of the load balancer.

Click **Next**.

5. On the Assign Servers to Clusters screen, assign both the servers that have been created to the cluster that you have just created.
6. Step through the Coherence Clusters screen.
7. On the Machines screen you can add machines in the cluster, and allocate managed servers to them. The AdminServer does not need to be allocated to a machine since it is started manually via the `startWebLogic.sh` script and is not managed by the Node Manager.
8. Step through the remaining screens and select **Create** to create the cluster.

4.6.2 Starting EDQ in a Cluster

Once the domain has been created, start the Admin Server by running the following command:

```
FMW_HOME/user_projects/domains/edq_domain/bin/startWebLogic.sh
```

Then start the Node Manager with the command:

```
FMW_HOME/user_projects/domains/edq_domain/bin/startNodeManager.sh
```

You can then log onto the Admin Console through a browser and start both of the Managed servers that were created during the configuration steps.

You can access the servers as follows:

- You can access any of the servers' launchpads. Connect to one of the servers and start the Director application. Or,
- You can connect via a load balancer in front of the cluster, and will be connected to whichever managed server the load balancer picks.

To display a dialog indicating which server the GUI is connected to, right-click the server in the project browser and select **Server Information**.

4.6.3 Enabling JMX API and Command Line for HA Clusters

The EDQ high availability deployment templates come with the internal JMX server disabled. This is to prevent port clashes between multiple managed servers which are running on the same host. Since the internal JMX server is not running the various command line tools which use this API, it will be unable to access the server. In particular the command line tools for starting and cancelling jobs will not work.

To enable these tools the internal JMX server must be re-enabled. This is achieved by editing the `director.properties` file in the `oedq.local.home` directory. Add the following line to the `director.properties` file:

```
management.localserver = "true"
```

This line enables the JMX server. By default the JMX server runs on port 8090. If multiple EDQ managed servers will run on the same host then you must configure each server to use a different port number.

Provided the EDQ managed servers are named with a number at the end, such as 'edq_server1', 'edq_server2', etc. (this is the default naming scheme), adding the line:

```
management.port = 8090 + servernum - 1
```

to the same `director.properties` file will cause the first EDQ managed server to run the JMX server on port 8090, the second server to run it on port 8092 and so on.

4.6.4 Landing Area

The landing area is a feature that allows EDQ to read and write data to server file system. If the landing area is to be used in a cluster then some consideration needs to be given as to how the landing area is shared amongst the managed servers. By default the landing area is located in the `oedq.local.home` area. If this is shared amongst the various hosts supporting EDQ managed servers then the landing area will continue to work as in a non-clustered system.

If the landing area file system is not shared amongst the managed servers, but is required for use then a number of options are available:

- The location of the landing area can be changed using the `landingarea` property in the `director.properties` file to a location that is shared amongst the hosts running EDQ managed servers.
- Since an EDQ job runs all its tasks and processes on the same managed server, any files consumed or generated by the job are written to its local landing area. External tasks can be added to the job to transfer any incoming or outgoing files to an appropriate shared location.
- If an EDQ job is consuming external files then these could be copied to all managed servers before the job is started.
- If an EDQ job is generating files for consumption by further EDQ jobs then the landing area can be synchronized across the various managed server between the various job run by use of an external tool such as `rsync`.

If files are being generated and consumed within the same job then a shared landing area may not be necessary. This is because the entire job will run on the same managed server and so access the local landing area.

Configuring Enterprise Data Quality with Apache Tomcat

This chapter describes how to configure EDQ and the database to operate with a Tomcat Application Server and perform the necessary additional configuration steps.

Note: These instructions apply to Apache Tomcat environments only. If you are using Oracle WebLogic Server, you must follow the directions in [Chapter 4, "Configuring Enterprise Data Quality with Oracle WebLogic Server."](#)

This chapter includes the following sections:

- [Prerequisites for these Procedures](#)
- [Creating the EDQ Database Objects](#)
- [Configuring Tomcat Application Server](#)
- [Creating Directories for Use With Tomcat Application Server](#)
- [Configuring EDQ to work with Tomcat Application Server](#)
- [Verifying EDQ Functional Packs](#)
- [Deploying the EDQ Application on a Tomcat Application Server](#)

5.1 Prerequisites for these Procedures

Before performing the procedures in this section, you must first read and satisfy the steps in:

- [Chapter 1, "Planning an Enterprise Data Quality Installation"](#)
- [Chapter 2, "Installing the Required External Software Components"](#)
- [Chapter 3, "Installing Enterprise Data Quality"](#)

5.2 Creating the EDQ Database Objects

EDQ requires three database accounts and three schemas. The first schema contains configuration data, the second contains results data, and the third contains staging data. These objects must be created in the database that you installed as the EDQ repository in [Chapter 2, "Installing the Required External Software Components."](#)

5.2.1 Creating an EDQ Repository in an Oracle Database

This procedure configures an Oracle database as the repository for EDQ.

1. Create three database accounts named `edqconfig`, `edqresults`, and `edqstaging` or other names of your choosing. When selecting names, consider that one schema will be the configuration schema and the other will be the results schema.
2. Grant the following privileges to each of these accounts:
 - `CREATE SESSION` (normally granted by `CONNECT`)
 - `CREATE INDEXTYPE` (normally granted by `RESOURCE`)
 - `CREATE SEQUENCE` (normally granted by `RESOURCE`)
 - `CREATE TABLE` (normally granted by `RESOURCE`)
 - `CREATE TYPE` (normally granted by `RESOURCE`)
 - `CREATE INDEX`
 - `DROP TABLE`
 - `DROP INDEX`
 - `UNLIMITED TABLESPACE` (recommended to allow the auto-extension of tablespace)

5.2.2 Creating an EDQ Repository in a PostgreSQL Database

This procedure configures a PostgreSQL database as the repository for EDQ.

1. Create three EDQ users named `edqconfig`, `edqresults`, and `edqstaging` or other names of your choosing.
2. Grant the database administrator (DBA) privilege to the users that you created.
3. Create the following schemas:
 - A schema named `edqconfig` or another name of your choosing. This schema must be owned by the first user you created and will contain the configuration data.
 - A schema named `edqresults` or another name of your choosing. This schema must be owned by the second user that you created and will contain the results data.
 - A schema named `edqstaging` or another name of your choosing. This schema must be owned by the third user that you created and will contain the staging data.
4. Test your database configuration by logging out, then back in, using the credentials for the new EDQ users.

5.3 Configuring Tomcat Application Server

Create the necessary Oracle Application Development Framework (ADF) shared library:

1. Create a `lib.adf` ADF shared directory in your Tomcat directory.
2. Unzip the `adf-essentials.zip` file, delivered with the EDQ product, into the `lib.adf` directory. The JAR files in this zip file must be in the `lib.adf` directory, not the `adf-essentials` directory that it unzips to by default.

3. Configure the Tomcat `common.loader` property (typically located in the `catalina.properties` file) to load Oracle ADF from their respective shared library paths. For example:

```
common.loader=${catalina.base}/lib,${catalina.base}/lib/*.jar,
${catalina.home}/lib,${catalina.home}/lib/*.jar,
${catalina.home}/lib.adf/*.jar
```

5.4 Creating Directories for Use With Tomcat Application Server

Create two empty directories to contain your EDQ configuration files. For example:

On Linux and UNIX operating systems:

```
/opt/edqconfig/oedqhome
```

```
/opt/edqconfig/oedqlocalhome
```

On the Windows operating system:

```
C:\edqconfig\oedqhome
```

```
C:\edqconfig\oedqlocalhome
```

The first directory (`oedqhome`) is the base configuration directory and will contain the configuration files that should not be changed post-installation. The second directory (`oedqlocalhome`) is the local configuration directory and will contain any custom settings that you create. Ensure that your application server user has read and write access to the two directories you create. For more information about these directories, see [Section 1.4.2, "EDQ Directory Requirements."](#)

5.5 Configuring EDQ to work with Tomcat Application Server

Follow these steps to run the EDQ Configuration Application (`configapp`) to configure EDQ to work with Apache Tomcat and to populate the repository schemas with the required EDQ objects.

Note: These instructions assume you have created the required database schemas as directed in [Section 2, "Installing the Required External Software Components."](#)

To start the EDQ Configuration Application, follow these steps:

1. Log in to the system as your EDQ installation user.
2. Go to your `FMW_HOME/edq/oracle.edq` directory.
3. Start the Configuration Application by entering the following command:

```
java -jar configapp.jar
```

The EDQ Configuration Application is displayed.

Use [Table 5-1, "Running the EDQ Configuration Application"](#) to configure EDQ:

Table 5-1 *Running the EDQ Configuration Application*

Screen	Action to Perform
EDQ Configuration Application	Click Begin to proceed with the configuration.
Completion	You may cancel the installation at any time by clicking Cancel .

Table 5–1 (Cont.) Running the EDQ Configuration Application

Screen	Action to Perform
Configuration Directory	Specify the two empty directories that you created in Section 5.4 to contain your EDQ configuration files. Click Next to continue.
Functional Packs	Select the functional packs you want and are licensed to use. Click Next to continue.
Configure the Configuration Schema	From the Selection Type list, select the correct option for the <i>configuration</i> schema that you created in Section 5.2 . <ul style="list-style-type: none"> Do not use the JNDI Connection option. Enter the user name and password for the configuration schema. Select the option for the type of database, either Oracle or Postgres . Enter the host name of the system that contains the database, its port number, and the database unique id. (For Oracle the unique id is the <i>SID</i> not the service name. If you want to use a pluggable database, which requires a service name, you must use a raw JDBC URL or JNDI source.) Ensure that you configured the connection to the database correctly by clicking Test . Click Next to continue.
Configure the Results Schema	From the Selection Type list, select the correct option for the <i>results</i> schema that you created in Section 5.2 . <ul style="list-style-type: none"> Do not use the JNDI Connection option. For PostgreSQL, you may want to use the Properties option. Enter the user name and password for the results schema. Select the option for the type of database, either Oracle or Postgres . Enter the host name of the system that contains the database, its port number, and the database unique id. (For Oracle the unique id is the <i>SID</i> not the service name. If you want to use a pluggable database, which requires a service name, you must use a raw JDBC URL or JNDI source.) Ensure that you have configured the connection to the database correctly by clicking Test . Click Next to continue.
Selected Options	Review the summary of your selections. Use Back to modify any of your selections. Click Finish to complete the configuration and continue.
EDQ Configuration Application Completion	Click Done to exit the installation program.

5.6 Verifying EDQ Functional Packs

You can verify which EDQ functional packs have been installed by using the EDQ Configuration Application as described in [Section 5.5, "Configuring EDQ to work with Tomcat Application Server."](#) Additionally, you should be aware of the following:

- All Functional Packs are needed if you want to use the full capabilities of the EDQ Customer Data Services Pack on your EDQ server, such as for integration with Siebel Customer Relationship Management or Universal Customer Master.
- All Functional Packs are needed if you want to install Oracle Watchlist Screening on your EDQ server.

5.7 Deploying the EDQ Application on a Tomcat Application Server

The EDQ application must be manually deployed on Tomcat Application Servers.

1. Stop the application server.
2. Deploy the `edq.war` file on your application server.

See the *Tomcat Web Application Deployment* web page at

<http://tomcat.apache.org/tomcat-8.0-doc/deployer-howto.html>

3. To link EDQ to the base and local configuration directories that you created in [Section 5.4](#), you can do one of the following:
 - On any of the EDQ-supported platforms, you can use a Java property named `-Dedq.config.path` that specifies the paths to those directories. The syntax is as follows, where `path_to_base_config` is the path to the base configuration directory and `path_to_local_config` is the path to the local configuration directory.

Linux and UNIX (colon separates the paths):

```
-Dedq.config.path=path_to_base_config:path_to_local_config
-Doracle.mds.cache=simple
```

Windows (semi-colon separates the paths):

```
-Dedq.config.path=path_to_base_config;path_to_local_config
-Doracle.mds.cache=simple
```

- On Windows, you can append an environment variable named `EDQ_CONFIG_PATH` to your `JAVA_OPTS` environment variable, as shown in the following, where `path_to_base_config` is the path to the base configuration directory and `path_to_local_config` is the path to the local configuration directory.

```
EDQ_CONFIG_PATH="path_to_base_config;path_to_local_config"
```

For more information about setting Java parameters, see the Apache Tomcat Documentation website at

<http://tomcat.apache.org/>

4. Restart your application server service and ensure that `edq.war` is successfully deployed.

Setting Server Parameters to Support Enterprise Data Quality

This chapter contains recommended Managed Server parameter settings for the framework that supports EDQ. These settings apply to both Oracle WebLogic Server and Apache Tomcat application server environments.

This chapter contains the following topics:

- [How to Set the Server Parameters](#)
- [Recommended Parameter Values](#)
- [Example of Parameter Settings](#)

6.1 How to Set the Server Parameters

This section describes how to set the recommended server parameters for WebLogic Server and Tomcat domains. Follow the guidelines that are appropriate for your EDQ installation.

6.1.1 How to Set Server Parameters in a WebLogic Server Domain

When you start a Managed Server in a WebLogic Server domain, the startup passes a set of configuration parameters to initialize a Java Virtual Machine (JVM). These parameters include, but are not limited to, the server's JVM settings and the path to the server configuration directories.

The method of setting the required server configuration parameters for EDQ varies, depending on the method that you use to start the Managed Servers. The following points explain the differences:

- When starting an EDQ Managed Server from the WebLogic Server Administration Console, the configuration settings for the Managed Server are pre-defined as server startup arguments, because they were set by the domain Configuration Wizard (see [Section 4.3](#)). When starting the server, these arguments are applied by default, but you can change them to the recommended settings described in [Section 6.2](#).
- When starting an EDQ Managed Server by using a startup script, the server startup arguments that are defined in the WebLogic Server Administration Console are ignored. However, when the domain Configuration Wizard was run, it created an `edq_server1` Managed Server that contains default JVM settings. These settings are stored in the `setStartupEnv.cmd` (Windows) or `setStartupEnv.sh` (Linux) script. You can edit this script to add or change the name of the Managed Server and to apply the recommended JVM settings

described in [Section 6.2](#).

6.1.2 How to Set Server Parameters in a Tomcat Domain

Consult the Tomcat documentation for your version of the application server. Documentation for all versions of Tomcat are at

<http://tomcat.apache.org>

6.2 Recommended Parameter Values

This section describes the recommended server parameter settings.

6.2.1 Setting JVM Parameters

You must set the following parameters for the Java Virtual Machine (JVM) that runs the EDQ environment. Most of these parameters apply to the HotSpot JVM, which is the central JVM of the Java Development Kit (JVM) from Oracle. This is the preferred JVM because it yields the best performance for EDQ.

- Maximum heap memory, `-Xmx`, should normally be set to approximately 50% of available physical memory on the host server. For example, on a server with 32GB of RAM, start by setting it to 16GB with `-Xmx16384m`. Depending on the relative usage of memory between the JVM and native processing, you may need to be adjust this setting. This parameter is common to most JVMs.
- Reserved Code Cache size should be set to 128m, `-XX:ReservedCodeCacheSize=128m`, unless otherwise advised. This parameter is specific to HotSpot JVM.

6.2.2 Setting Other Parameters

The following additional parameter setting is recommended.

- Headless mode should always be set to true on UNIX systems, `-Djava.awt.headless=true`. This parameter is specific to the Java Abstract Window Toolkit (AWT) library. Setting this is recommended because some of the graphic drawing routines (such as around graphs) will attempt to connect to the display device on a client to size fonts and other attributes. If the machine does not have a display device, this may interfere with operation. The `headless=true` setting directs the graphics library to operate without connecting to a display.

6.3 Example of Parameter Settings

The following is an example of the full set of recommended parameter options for a typical UNIX or Linux server with 32GB of RAM:

```
-Xmx16384m  
-XX:ReservedCodeCacheSize=128m  
-Djava.awt.headless=true
```

Next Steps After Configuring Enterprise Data Quality

This chapter describes how to start using EDQ and is intended to help you become familiar with the main components of EDQ.

This chapter includes the following sections:

- [Logging Into EDQ](#)
- [The EDQ Launchpad](#)

7.1 Logging Into EDQ

You can access the EDQ Launchpad and client applications by starting a supported browser and enter the following URL:

```
http://server_name:port_number/edq
```

where *server_name* is the name of the server onto which you installed EDQ and *port_number* is the HTTP or HTTPS port that your application server is running against (8001 on WebLogic Server by default, 8080 on Tomcat). If you deployed the application server to run against a different port, you should use that port number.

Enter your login credentials for the appropriate application server:

WebLogic Server Login

Use the WebLogic Server user that you created in [Section 4.3.2, "Navigating the Domain Configuration Wizard Screens."](#)

Tomcat Login

Use the default EDQ administrator account `dnadmin` with the password `dnadmin`. You are prompted to change the password the first time you log in. This password must meet the default security standards.

7.2 The EDQ Launchpad

The Launchpad provides access to the EDQ client applications, services, and system management. The following may be displayed on the Launchpad when EDQ is installed:

Director	Starts the Director client application, which is the main configuration application.
-----------------	--------------------------------------------------------------------------------------

Server Console	Allows you to perform the following tasks: <ul style="list-style-type: none"> ■ Schedule jobs ■ View current tasks ■ View the event log ■ View job results
Dashboard	Starts the Dashboard, where published data quality metrics are displayed.
Match Review	Starts the Match Review application, which allows users to view an overview of the reviews assigned to them and to launch the review application.
Case Management	Case Management is an application designed to support the manual investigation of results from data quality processes. It is also used as the main investigation application in Oracle Watchlist Screening, for both batch and real time screening results. Using Case Management, privileged users can manage and review matching results using highly configurable workflows with a comprehensive audit history of all investigation work.
Case Management Administration	Case Management Administration provides the following areas of functionality: <ul style="list-style-type: none"> ■ Workflow Administration—allows you to create, edit, copy, import, export and delete workflows. ■ Case Source Administration—allows you to import, export and delete case source definitions. ■ Permission Administration—allows you to define and manage the data that users have permission to access.
Administration	This application allows a sufficiently privileged user to configure EDQ users, permission groups, password and security rules, extensions, functional packs, the applications that appear on the launchpad, and the ability to view and monitor sessions.
Web Services	Displays details of the web services configured in the EDQ server, and allows access to the web service tester.
Login	An explicit Login button in the upper right corner. Upon logging in the full list of pages available to the user is displayed. Once the user is logged in they have access to a user menu by clicking on their user name in the top right of the screen. This provides access to log out and change password functionality (change password is only shown if using an internal realm).
Change Password	Allows users to change their passwords. A user must log in with the existing password and then provide and confirm the new password. This option is not available when using a WebLogic Server.
Help	Allows you to retrieve information about how to use each EDQ application.

An Administrator can reconfigure the Launchpad to define the user applications and links that are displayed to users by using the Administration pages. For more information, see the *Oracle Enterprise Data Quality Online Help*.

Upgrading Enterprise Data Quality

To upgrade EDQ running on WebLogic Servers or Tomcat Servers from Release 11g or Release 12c (12.1.3) to Release 12c (12.2.1) of EDQ, please see *Oracle Fusion Middleware Oracle Enterprise Data Quality Upgrade Guide*.

Note: You must review Oracle Fusion Middleware upgrade documentation before referring to the Oracle Fusion Middleware Oracle Enterprise Data Quality Upgrade Guide.

Read *Planning an Upgrade of Oracle Fusion Middleware* for additional guidelines for preparing to upgrade to Oracle Fusion Middleware Release 12c (12.2.1). This documentation also includes descriptions of terminology changes that you must understand to move forward to a 12c environment.

Removing Enterprise Data Quality from a System

This chapter describes how to remove EDQ from Linux, UNIX, and Windows.

This chapter includes the following sections:

- [Removing EDQ from a Linux or UNIX System](#)
- [Removing EDQ from a Windows System](#)

Note: These procedures remove the EDQ files that were installed by Oracle Universal Installer (OUI). They do not remove the EDQ application from Oracle WebLogic Server. Follow the instructions in "Removing Your Domain and Application Data" in *Oracle® Fusion Middleware Installing and Configuring Oracle WebLogic Server and Coherence*.

9.1 Removing EDQ from a Linux or UNIX System

To remove EDQ from a Linux or UNIX system, follow these steps:

1. Log in to the target system as your EDQ installation user.
2. Go to the directory, `EDQ_HOME/oui/bin`.

Note: Ensure that all files in the directory have executable permissions for the user you logged in with before continuing.

3. Run the following command to start the `deinstall.sh` program:

```
./deinstall.sh
```

The EDQ deinstallation program is displayed.

4. Select Enterprise Data Quality 12.2.1.0.0 as the distribution to remove, and then click **Uninstall**.
5. Navigate through the screens and then click **Finish** to complete the removal of EDQ from the system.

9.2 Removing EDQ from a Windows System

To remove EDQ from a Windows system, follow these steps:

1. Log in to the target Windows system as your EDQ installation user.
2. Locate the MS-DOS Command Prompt (`cmd.exe`), right-click on it, and then select **Run as administrator**.
3. Go to the installation directory `EDQ_HOME\oui\bin`.
4. Enter the following command:

```
deinstall.cmd -deinstall
```

The EDQ deinstallation program is displayed.
5. Select Enterprise Data Quality 12.2.1.0.0 as the distribution to remove, and then click **Uninstall**.
6. Navigate through the screens and then click **Finish** to complete the removal of EDQ from the system.