

Oracle® FMW

Release Notes for Enterprise Data Quality



12c (12.2.1.4.4)

F36176-08

March 2024

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

ORACLE®

Oracle FMW Release Notes for Enterprise Data Quality, 12c (12.2.1.4.4)

F36176-08

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

Related Documents	vi
Documentation Accessibility	vi
Conventions	vi

1 Oracle Enterprise Data Quality Release Notes

1.1 Release 12.2.1.4.4	1-1
1.1.1 New Features and Improvements	1-1
1.1.1.1 Web Push Notification Support	1-1
1.1.1.2 OCI Queue Service Support for Realtime Processing	1-1
1.1.1.3 Case Management Script Library	1-2
1.1.1.4 Redwood UI for Enterprise Data Quality	1-2
1.1.1.5 REST APIs for Packaging and Importing Configuration	1-2
1.1.1.6 Automatic Update of Oracle Schema Password and Wallet Refresh	1-2
1.1.1.7 Folder for Local Web Content	1-2
1.1.1.8 Support for Multi-part File Uploads	1-3
1.1.1.9 Support for Apache Derby Database for Schema Configuration	1-3
1.1.1.10 REST APIs for User and Launchpad Administration	1-3
1.1.1.11 Limited Support for Global Web Services	1-3
1.1.1.12 Comparison Counts in Realtime Match Processing	1-3
1.1.1.13 Enhanced Encryption Support	1-4
1.1.2 Applying a Bundle Patch for Oracle Enterprise Data Quality	1-5
1.1.3 Issues Resolved for 12.2.1.4.4	1-5
1.1.4 Known Issues and Workarounds	1-12
1.2 Release 12.2.1.4.3	1-12
1.2.1 New Features and Improvements	1-12
1.2.1.1 Support for Long VARCHAR Columns in EDQ Results Schema	1-12
1.2.1.2 Enhanced Redirect URI Selection for OpenID Authentication	1-13
1.2.1.3 Case Management REST API Enhancements	1-13
1.2.1.4 Support for Amazon Simple Queue Service	1-13
1.2.1.5 Support for User Authentication using Amazon Cognito User Pools	1-13
1.2.1.6 Java Flight Recorder tool to Monitor Web Service Requests	1-13

1.2.1.7	Support for Apache Tomcat Version 10	1-13
1.2.1.8	Adding Additional JMS Drivers in <messengerconfig>	1-14
1.2.1.9	Apache Tomcat Clustering Support for EDQ	1-14
1.2.1.10	Connectivity to Apache Parquet Data Stores	1-14
1.2.1.11	Configuring Additional HTTP Response Headers	1-15
1.2.1.12	Tool to Update Database Passwords in director.properties	1-15
1.2.1.13	EDQ Application Launchers for Mac Operating Systems	1-16
1.2.1.14	Support for Downloading Multiple Files from Cloud Storage	1-16
1.2.1.15	Azure AD Authentication Support for Autonomous Database	1-17
1.2.1.16	Utility to Determine the MIME type of Uploaded Files	1-17
1.2.2	Applying a Bundle Patch for Oracle Enterprise Data Quality	1-17
1.2.3	Issues Resolved for 12.2.1.4.3	1-18
1.2.4	Known Issues and Workarounds	1-21
1.2.4.1	Using Tomcat with Java 8 Causes EDQ to Fail	1-21
1.3	Release 12.2.1.4.2	1-21
1.3.1	New Features and Improvements	1-21
1.3.1.1	Support for User Authentication using Azure Active Directory	1-22
1.3.1.2	Support for User Authentication using Oracle Identity Cloud Service	1-22
1.3.1.3	Enhancements to the Apache Hive™ Connector	1-22
1.3.1.4	EDQ Application Launchers for Supported Client Operating Systems	1-22
1.3.1.5	Connectivity to Snowflake	1-22
1.3.1.6	New Case Management API	1-23
1.3.1.7	Support for Java versions 11 and 17	1-23
1.3.1.8	Improved Reporting for Case Management Filter Execution	1-23
1.3.2	Applying a Bundle Patch for Oracle Enterprise Data Quality	1-23
1.3.3	Issues Resolved for 12.2.1.4.2	1-23
1.3.4	Known Issues and Workarounds	1-28
1.3.4.1	Using Tomcat with Java 17 Causes Snowflake Snapshots to Fail	1-28
1.4	Release 12.2.1.4.1	1-29
1.4.1	New Features and Improvements	1-29
1.4.1.1	Stored Credentials	1-29
1.4.1.2	Extended Limit of Oracle VARCHAR Columns in EDQ Results Schema	1-29
1.4.1.3	Data Store for Autonomous Databases	1-29
1.4.1.4	Connectivity to Apache Kafka	1-30
1.4.1.5	Connectivity to Apache Hive™ using Kerberos Authentication	1-30
1.4.1.6	Connectivity to Apache Avro™ Data Stores	1-30
1.4.1.7	Connectivity to AWS Redshift	1-30
1.4.1.8	Connectivity to JSON Data Stores	1-30
1.4.1.9	Connectivity to Oracle Service Cloud	1-30
1.4.1.10	Support for Case Management Filter and Report Execution by using SQL and Oracle Text	1-31

1.4.1.11	Support for Reconnect Retries in JMS Connection	1-31
1.4.2	Applying a Bundle Patch for Oracle Enterprise Data Quality	1-31
1.4.3	Removal Notice	1-32
1.4.4	Issues Resolved	1-32
1.4.5	Known Issues and Workarounds	1-35
1.4.5.1	Opening Server Package File in Apple OSX Causes Error	1-35
1.5	Release 12.2.1.4.0	1-35
1.5.1	New Features and Improvements	1-35
1.5.1.1	Call External Web Service Processor	1-35
1.5.1.2	Cassandra Data Store	1-36
1.5.1.3	Multiple Value Reference Data Editing	1-36
1.5.1.4	New REST API Documentation and Testing capabilities	1-36
1.5.2	Issues Resolved	1-36
1.5.3	Known Issues and Workarounds	1-39
1.5.3.1	Opening Server Package File in Apple OSX Causes Error	1-39
1.5.4	Upgrade Considerations	1-39
1.5.4.1	Upgrades in an Oracle WebLogic Server Environment	1-39
1.5.4.2	Upgrades in an Apache Tomcat Environment	1-40

A Switching to SQL and Oracle Text for Case Management Filters and Reports

A.1	Working of Case Management Filters	A-1
A.2	Understanding Oracle Text Expressions	A-1
A.3	Using Oracle Text Options	A-2
A.4	Updating Schema	A-2
A.5	Creating Helper Index and Column	A-2
A.6	Creating Text Indices	A-3
A.7	Populating JSON	A-3
A.8	Enabling SQL and Oracle Text Usage for Filtering and Reports	A-4
A.9	Additional Considerations	A-4

B Manually Patching Oracle Enterprise Data Quality

B.1	Installing a Patch on Weblogic	B-1
B.2	Installing a Patch on Tomcat	B-3

C Reporting for Case Management Filter Execution

Preface

Related Documents

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

EDQ Documentation Library

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 user applications of Oracle Enterprise Data Quality. 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**.

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.

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.

Convention	Meaning
<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.

1

Oracle Enterprise Data Quality Release Notes

Oracle recommends you review this document before installing or working with the product.

It has the following sections:

- [Release 12.2.1.4.4](#)
- [Release 12.2.1.4.3](#)
- [Release 12.2.1.4.2](#)
- [Release 12.2.1.4.1](#)
- [Release 12.2.1.4.0](#)

1.1 Release 12.2.1.4.4

This document contains release information for Oracle Enterprise Data Quality (EDQ) (12.2.1.4.4) and includes the following sections:

- [New Features and Improvements](#)
- [Applying a Bundle Patch for Oracle Enterprise Data Quality](#)
- [Issues Resolved for 12.2.1.4.4](#)
- [Known Issues and Workarounds](#)

1.1.1 New Features and Improvements

This release introduces the following enhancements:

1.1.1.1 Web Push Notification Support

EDQ 12.2.1.4.4 and later include web notifications support.

These notifications are generated by triggers that can report on job start/stop/error and other events. Notifications can also be used for infrequent events that may demand immediate attention from users. Note that these notifications are not meant to replace existing schemes like emails for more frequent events such as case management assignment. For more information about enabling web push notifications, see [Configuring EDQ Web Push Notifications](#).

1.1.1.2 OCI Queue Service Support for Realtime Processing

EDQ 12.2.1.4.4 and later support realtime integration with OCI queues.

OCI has introduced a new queue service, which is a much better fit for realtime processing as compared to OCI streams. EDQ realtime provider and consumer buckets can now be configured to use OCI queues for reading and publishing. For more information, see [Using Oracle Cloud Infrastructure \(OCI\) Queue with EDQ](#).

1.1.1.3 Case Management Script Library

EDQ 12.2.1.4.4 introduces a case management script library, which can be used in EDQ triggers and script processors, to access Case Management data and a number of functions in a supported manner.

This avoids issues some customers have faced with attempts to use internal interfaces to Case Management data.

For more information, see [Using Case Management Scripting](#).

1.1.1.4 Redwood UI for Enterprise Data Quality

EDQ now uses the Oracle Redwood theme to make for a standardized user experience across Oracle products.

The Launchpad and help page user interfaces now display in the Redwood color theme.

1.1.1.5 REST APIs for Packaging and Importing Configuration

EDQ 12.2.1.4.4 includes new APIs for packaging configuration (especially projects) and for importing configuration to facilitate automated transfer of configuration between EDQ environments.

These configuration objects include Projects, Global Reference Data, Global Data Stores, Global Published Processors, Stored Credentials, Case Sources, Case Workflows, and Case Permissions. For more information, see [Using REST APIs for Importing and Exporting Configuration](#).

1.1.1.6 Automatic Update of Oracle Schema Password and Wallet Refresh

The default profile on many Oracle database installations sets passwords to expire automatically.

EDQ 12.2.1.4.4 now includes the capability to check the configuration and results schemas periodically and issue a warning if the password will expire within a configurable interval. For EDQ running on Tomcat, you can also configure automatic reset of the schema password.

For more information, see [Configuring Schema Password Expiry Warnings and Wallet Refresh](#).

1.1.1.7 Folder for Local Web Content

EDQ 12.2.1.4.4 adds a new "local" web folder to store additional static web content to an EDQ application installation, for example, to support the display of a custom image in a web notification, or in an extended user application.

The URL for files in the local area is:

```
http://server:port/edq/local/
```

The default location for the directory containing local content is `localcontent` in the EDQ "local home" configuration directory. You can set an alternative location by setting **localcontent.directory** in *director.properties*.

You can add files to the local content directory using any of these mechanisms:

- Normal OS copy operations
- the built-in EDQ sshd server
- POST or PUT requests to `http://server:port/edq/local/path`

For more information, see *Using the Local Web Content Directory*.

1.1.1.8 Support for Multi-part File Uploads

The EDQ Upload File task normally sends the data in a single HTTP PUT or POST call. Some cloud providers impose a limit on the size of the file that can be uploaded in a single request.

In EDQ 12.2.1.4.4 the Upload File task is enhanced to support multipart uploads. In a multipart upload the file is divided into parts that are sent in separate HTTP calls. Multiple parts can be uploaded concurrently. Once all the parts have been uploaded, the parts are assembled into a single file using a multipart commit call.

1.1.1.9 Support for Apache Derby Database for Schema Configuration

EDQ 12.2.1.4.4 and later on Tomcat support the use of an embedded Apache Derby database for the configuration and results schemas.

1.1.1.10 REST APIs for User and Launchpad Administration

EDQ 12.2.1.4.4 onwards you can use REST APIs to manage groups, permissions, external group mappings, users, and launchpad applications.

For more information, see *Using REST APIs for User Management and Launchpad Administration*.

1.1.1.11 Limited Support for Global Web Services

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 restores limited support for global webservices. You can define provider (reader) and consumer (writer) interfaces using real-time bucket files. Script decoding is supported 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 that SOAP is not supported. See *Using Scripted Global Web Services with EDQ*.

1.1.1.12 Comparison Counts in Realtime Match Processing

The count of comparisons performed in a Match processor is useful for diagnosing performance issues. EDQ 12.2.1.4.4 onward you can determine the comparison count in realtime mode.

The count is exposed as a message tag named `match.comparisons` and in a custom Java Flight Recording (JFR) event.

In REST/JSON web service calls, message tags are added as extra HTTP headers in the output:

```
X-EDQ-MESSAGE-TAG: match.comparisons=466
```

In SOAP or REST/XML web service calls, message tags are added as attributes in the top level XML element:

```
<?xml version="1.0" encoding="UTF-8" standalone="no"?>
<soap:Envelope xmlns:soap="http://schemas.xmlsoap.org/soap/envelope/">
  <soap:Body>
    <dn:response xmlns:dn="http://www.datanomic.com/ws"
match.comparisons="466">
  ...
```

Alternatively, the Get Message Header Number processor can be used to extract the tag into a process attribute. This may be useful if the process chain contains more than one Match processor. There is a single tag value only and the count from a later processor would overwrite earlier values.

EDQ running with Java 11 or later generates a JFR event for each realtime message processed. Events are in the EDQ > Match processing > Realtime category and contain these attributes:

Attribute	Description
Start Time	Timestamp of event start.
Duration	Event duration.
End Time	Timestamp of event end.
Event Thread	Name of thread in which the event originated.
comparisons	Match comparison count.
ECID	Unique ID of request. In WebLogic this is generated by the diagnostic framework automatically. In Tomcat this is a random UUID.
process	Process name.
processor	Processor name.

1.1.1.13 Enhanced Encryption Support

The encryption framework EDQ 12.1.1.4.4 now supports a pluggable module with implementations for OCI vaults and scripts.

The previous encryption mechanism for data source passwords and stored credentials did not support ready migration of database schemas from old to new systems. In addition to this, the cipher used ECB mode that generated the same result each time for given inputs. Now encryption and decryption is handled by a pluggable security module to handle encryption and, optionally, credentials (secrets) retrieval.

There may be cases where credentials (username and password) are stored in one or more additional external stores. For example internal encryption may be used but there is a requirement to extract credentials from OCI vaults or AWS secrets manager. For such case, the security module framework is extended to support additional credentials stores.

For secrets retrieval or encryption APIs that cannot readily be called from scripts, this release also allows you to implement modules using custom Java classes.

1.1.2 Applying a Bundle Patch for Oracle Enterprise Data Quality

A bundle patch is an official Oracle patch for Oracle Fusion Middleware components on baseline platforms. Each bundle patch includes the libraries and files that have been rebuilt to implement one or more fixes. All of the fixes in the bundle patch have been tested and are certified to work with one another.



Note:

Oracle recommends that you have the latest version of Opatch (version 13.9.2.0.0+ or later) from My Oracle Support. Opatch requires access to a valid Oracle Universal Installer (OUI) Inventory to apply patches.

Oracle Fusion Middleware 12.2.1 products are installed with OPatch NextGen 13.9.2.0.0 to apply interim patches. The Oracle patch mechanism (Opatch) is a Java-based utility that runs on all supported operating systems. Opatch requires installation of the Oracle Universal Installer.

For detailed instructions, refer to [Manually Patching Oracle Enterprise Data Quality](#).

1.1.3 Issues Resolved for 12.2.1.4.4

This section describes issues resolved in this release.

Table 1-1 Issues Resolved

Issue	Notes
XLSX files with repeated string data returns a zip bomb error when opened for reading	<p>Certain XLSX files that have a lot of repeated string data could provoke a zip bomb error when opened for reading:</p> <pre>java.io.IOException: Zip bomb detected! The file would exceed the max. ratio of compressed file size to the size of the expanded data.</pre> <p>The error could happen when the xlsx file is opened for snapshotting or when you try to export the file without the "always overwrite" option set on the data store.</p> <p>You can now use the Java system property <code>edq.xlsx.zipratio</code> to set the <code>ZipSecureFile.setMinInflateRatio</code> value. You can set it in <i>jvm.properties</i> or using options in the command line tool used to start the application server.</p> <p>Note that the ratio displayed in the error message is an estimate, which may change as more of the file is read. A number of iterations may be required to find the best value for a particular input file. Set the value to 0.0 to suppress all zip bomb checks.</p>
Exports to xls(x) files fail with a null pointer exception if the target worksheet is empty	<p>Exports to xls(x) files would fail with a null pointer exception if the target worksheet exists but was empty. This occurred even if the always overwrite option was enabled. The workaround for this issue was to delete the empty file before the export. This issue has been fixed in this release.</p>
Create permission is required to modify a web service	<p>A config API call to modify a web service failed if the user did not have the <code>add web service</code> permission. This issue has been fixed in this release. Users must have the <code>modify web service</code> permission to do this.</p>
User without modify group permission cannot access User Groups page	<p>User with Add Group permission could not access the User Groups page. The page would display a 403 error. The group edit page loaded the available permissions and the call for this insisted that the user have the modify group permission. This issue has been fixed in this release.</p>
Data store config API does not save or return description field	<p>The data store object used in the config API includes a description field. When you create a store using the API, the description is ignored. When data stores were retrieved by the API, the existing descriptions were not included. This issue has been fixed in this release.</p>

Table 1-1 (Cont.) Issues Resolved

Issue	Notes
Autonomous Database data stores do not refresh wallet files	<p>If you create an Autonomous DB data store using automatic wallet download, the wallet is fetched and stored only once when the first connection is made. The wallet is never downloaded again. If the wallet certificate expires, then subsequent connections would fail to load. This issue has been fixed in this release.</p> <p>Wallet files that are older than 30 days are refreshed at a configurable interval. You can specify the refresh interval by setting the following:</p> <pre>autonomousdb.wallet.refresh.interval = N</pre> <p>where N can be suffixed with d (for days), h (for hours), m (for minutes) or s (for seconds). If no suffix is specified, the suffix is assumed to be d by default.</p>
Error when viewing alert if it contains an attachment with invalid name	<p>For cases and alerts that include attachments, Case Management tries to display the correct icon corresponding to the attachment. For cases and alerts that contained an attachment whose file name was not valid according to the operating system of the client, an attempt to view such cases or alerts would fail with a <code>java.io.file.InvalidPath</code> exception.</p> <p>This issue has been fixed in this release. Now such exceptions when caught are ignored.</p>
Error displayed if a change is made to case or alert that makes it unavailable to user	<p>In case management a case or alert could change such that it is no longer available to a user. For example,</p> <ol style="list-style-type: none"> 1. If the case is assigned to a different user and the current user does not have the "View cases assigned to others" permission. 2. If the case or alert has transitioned and a permission is added to the case or alert. <p>In such cases, an error message is displayed before the case or alert details are blanked out. The user could proceed to work on other tasks after the message asking for confirmation is displayed and accepted.</p> <p>This issue has been fixed in this release. Cases or alerts that are unavailable after a change assignment are no longer displayed to that user.</p>
Add comment fails if user does not have restrict comment permission	<p>For users that do not have the <code>restrict comments</code> permission, the add comment dialog box on the client system did not show the Permission selection. This issue has been fixed in this release.</p>
Search result links in built-in help do not work in all browsers	<p>Links generated by a search in the embedded help did not work in all browsers. This issue has been fixed in this release.</p>
Remove RMI realtime messenger type	<p>The realtime framework supported an RMI messenger that used remote calls to submit requests. The RMI messenger does not work through a load balancer. This release onwards the RMI messenger is no longer included in EDQ installations.</p>

Table 1-1 (Cont.) Issues Resolved

Issue	Notes
Remove realtime bean and map decoding types	<p>The realtime bucket framework supported the JMS "object" (bean) and "map" message types. Neither of these message types are relevant to other realtime messengers such as SQS and Kafka. In addition to this, object messages use Java serialization, which can be a security risk. This release onwards support is removed for object and map message types. Text format, such as json, xml, and so on, is now the only supported message type.</p>
Realtime payload parsing scripts using Groovy do not work with Java 16 and later	<p>Scripts used to parse incoming message bodies construct 'Record' objects to contain attribute data. For example:</p> <pre data-bbox="769 737 1341 1014"> def extract(str, type, mtags) { def x = new XmlParser().parseText(str); def r = Record.new(); r.field1 = x.@field1; r.field2 = x.@field2.toDouble(); return [r]; } </pre> <p>Java 16 introduced the <code>java.lang.Record</code> class. Groovy interprets the <code>Record.new()</code> call as using <code>java.lang.Record</code> and hence the call fails at runtime with the error:</p> <pre data-bbox="769 1199 1320 1283"> No signature of method: static java.lang.Record.new() is applicable for argument types: () values: [] </pre> <p>This issue has been fixed in this release. Scripts can now use <code>newRecord()</code> instead of <code>Record.new()</code>.</p>
REST API to generate sysreport	<p>Support and customers find it difficult to generate sysreports using the existing mechanism. This issue has been fixed in this release. The following endpoint is now available in the system administration name space to generate a report in JSON format:</p> <pre data-bbox="769 1524 1344 1549"> GET http://server:port/edq/admin/sysreport </pre>
Web service MBeans are removed when project is overwritten during import	<p>When you import a process containing web services that overwrites the existing project, the monitoring MBeans for the web service are deleted. Such MBeans are removed when the old version of the web service is deleted. This issue has been fixed in this release.</p>
Support asynchronous calls in the HTTP script library	<p>The previous HTTP script library did not include an asynchronous calling mechanism. Customer currently using a Groovy script with an <code>ExecutorService</code> to run multiple web service requests in parallel could therefore not migrate such a script to JavaScript without significant performance issues as requests had to be sent serially. This issue has been fixed in this release.</p>

Table 1-1 (Cont.) Issues Resolved

Issue	Notes
Cache results of bearer auth access token user mapping	When web service calls are made using Bearer authentication with OAuth2 access tokens, the data in the token is mapped to a user identity for each call. The results should be cached to speed up subsequent calls. This issue has been fixed in this release.
Searches done using "modified by" filters for an extended attribute would fail with an SQL error	When running in Oracle Text/SQL mode, searches done using "modified by" filters for an extended attribute would fail with an SQL error: <code>java.sql.SQLException: ORA-01722: invalid number</code> This happened because the <code>custom_flagN_by</code> column in <code>dn_case</code> , which is a number was treated as a string when the SQL was created. This issue has been fixed in this release.
Deleting cluster in a match processor returns null pointer exception error	In a cluster configuration in a match processor, if you delete a cluster the client would fail with a <code>java.lang.NullPointerException</code> error. This issue has been fixed in this release.
Dashboard graphs in count mode show incorrect information	Dashboard graphs in count mode always show all available data, irrespective of the "view last" setting. This issue has been fixed in this release.
Users with read-only permission cannot view match and parse configuration	Users that only had read-only permission along with permissions to view configuration such as users in the default Data Stewards user group, could not open Match processors to view the configuration. In addition, such users were not able to see the set of match rules that are normally available when you click on the Match subprocessor. These issues also happened with parse configuration where such users could not open the subprocessors to view configuration. All these issues have been fixed in this release.
Remove Web Service Tester menu link in Director user interface	The Web Services menu in the Director user interface has options to launch the Web Service Tester and Dashboard. These links use the old ADF URLs, which do not work in EDQ 12.2.1.4.1 or later. These links are no longer available in EDQ 12.2.1.4.4.
Up and down arrows on the Launchpad configuration page did not work correctly	The up and down arrows on the Launchpad configuration page would sometimes fail to work correctly. This issue has been fixed in this release.
Add "basic" stored credentials type	For external calls such as download/upload/cews, which use "Basic" authentication, you needed to provide the user name and password. EDQ 12.2.1.4.4 introduces a Basic stored credentials type to centralize this. In addition, the password is stored in an encrypted format so that it cannot be read or modified.

Table 1-1 (Cont.) Issues Resolved

Issue	Notes
Java 11 bug can cause invalid XML to be stored within process model	<p>In EDQ 7.1 and earlier, the canvas layout for a process was stored in an embedded <code><guidata></code> element within the process model XML whereas in the recent versions the layout data is stored in a separate XML record. Processes that originated in EDQ 7.1 or earlier still have the embedded XML, even though it is not used, and the XML is preserved when the process is edited or copied.</p> <p>A bug in Java versions 9 to 13 causes additional white space to be added around <code><![CDATA</code> when the XML is output. This caused issues when you opened processes that were created in EDQ 7.1 or earlier. When such processes are opened and the copy is read back, the embedded <code><guidata></code> is parsed and the additional whitespace added as a result of the java bug becomes part of the embedded XML and the parsing fails because the <code><?xml . . .></code> instruction is no longer at the beginning of the XML.</p> <p>This issue has been fixed in this release. The <code><guidata></code> element is trimmed before parsing and is no longer stored when a process is edited.</p>
Cannot edit case and alert expiry time interval	<p>You could not edit the expiry time interval of a case or alert using the EDQ user interface UI or the REST API. Any changes made would be ignored. This issue has been fixed in this release.</p>
Case update version check fails to work	<p>Case and alert records have a "version" attribute, which is incremented each time the record is updated in the database. Every time the record is updated, the version is checked. The version check is intended to prevent users from updating a case/alert that has already been updated either by another user or by an API call after the user ran the filter to obtain the data.</p> <p>Security changes made in 12.2.1.4.3 to prevent unauthorized users from performing updates have broken the version check process. This has been fixed in this release.</p>
Signatures created with EC are not verified correctly	<p>JSON Web Token (JWT) access tokens used with OAuth2 Bearer authentication such as Azure, Cognito, IDCS contain a signature created with a private key owned by the issuer. These tokens can be signed either using RSA or EC algorithms.</p> <p>Signatures created with EC are not verified correctly. This is because the signature bytes are formed by concatenating two unsigned numbers whereas the Java ECDSA signature algorithm expects the numbers as an ASN.1 sequence. This issue has been fixed in this release.</p>

Table 1-1 (Cont.) Issues Resolved

Issue	Notes
Incorrect base64 encoding used with GCP tokens	The authentication mechanism using Google Cloud storage credentials requires the creation of a signed JSON Web Token (JWT) component that should be base64 encoded using the url-safe alphabet, and with trailing equal to (=) characters removed. However, the standard base64 alphabet was used, and the trailing equal to (=) characters were not removed. This issue has been fixed in this release.
Realtime eof predicates do not stop "run everywhere" job	The optional <code><eof></code> predicate in a realtime bucket definition will end input on the local reader, but does not stop a job that is running on all nodes in a cluster. This issue has been fixed in this release. A new shutdown attribute is added on the <code><eof></code> element. If set to true, the job is canceled with a shutdown.
Null pointer exception when using eof mechanism with SQS	If an SQS realtime provider is configured with an <code><eof></code> predicate, the process would fail with a null pointer exception when a message matching the predicate is received. To avoid this error use the following setting in the messenger properties <code>deletemode = completion</code>
Writing to existing XLSX file fails if Always Overwrite option is not selected	Exporting files to an XLSX data store that is created without the "Always overwrite" option would fail with an error like the following: "Attempting to write a row[0] in the range [0,10000] that is already written to disk." This issue has been fixed in this release.
Increase maximum length of column names to 128 for Oracle exports	Oracle database versions now support column names up to 128 bytes, which was limited to 30 bytes in older versions. EDQ Oracle DB connectors set the limit to 30. This limit should be increased if the database version and compatible setting are suitable. This issue has been fixed in this release.
Error thrown if a value longer than 8191 bytes was written to the last column during export to Oracle tables	When inserting data into an Oracle table with <code>string size = extended</code> , an ORA-24816 error would be thrown if a value longer than 8191 bytes was written to the last column. This issue has been fixed in this release.
Launchpad Sign In button functioned improperly	When you try to login to the Launchpad you are prompted to enter the username and password to sign in. The Sign In button would be disabled even after the password was entered. Pressing the ENTER key would enable the button, but the sign in action would not complete. This issue has been fixed in this release. The Sign In button is enabled when the password is entered and pressing the ENTER key allows the user to sign in.
Configuration Analysis Help content points to Director help page	The Help option within the Process > Configuration Reports would incorrectly display the help for the Director page. This issue has been fixed in this release.

Table 1-1 (Cont.) Issues Resolved

Issue	Notes
Notes can be added to jobs opened in server console	When a job is opened in server console it is in the read-only format. However, you could click the add note button and create a note on the canvas. This issue has been fixed in this release.
Projects imported from DXIs are restricted to Administrators group only	A project imported from a DXI were automatically given only to the Administrators group. This has been fixed in this release. Group names are now added to a project when exported and restored on import. To keep the original behavior on import, add the following to <code>director.properties</code> : <code>project.import.compatible = true</code>
Case Management REST API returns null pointer exception if comment is omitted during stage change update	If the Case Management update API is called with a state change specification that omits a comment, the update fails with a <code>NullPointerException</code> . This happens in cases where the transition does not require a comment or does not define any default comments. To avoid this issue, make sure that a comment is added for a case or alert.
Case Management REST API allows assignment to user without case permission	When the Case Management update API would assign a case or alert to a user, it would check whether the calling user has the permission, if any, assigned to the case instead checking the permission of the target user. This issue has been fixed.

1.1.4 Known Issues and Workarounds

This section details known issues in this release, and their workarounds.

1.2 Release 12.2.1.4.3

This document contains release information for Oracle Enterprise Data Quality (EDQ) (12.2.1.4.3) and includes the following sections:

- [New Features and Improvements](#)
- [Applying a Bundle Patch for Oracle Enterprise Data Quality](#)
- [Issues Resolved for 12.2.1.4.3](#)
- [Known Issues and Workarounds](#)

1.2.1 New Features and Improvements

This release introduces the following enhancements:

1.2.1.1 Support for Long VARCHAR Columns in EDQ Results Schema

EDQ 12.2.1.4.1 and 12.2.1.4.2 adapt to Oracle repository databases, which support extended string lengths to allow the capture of data up to 32767 bytes.

However, as all string columns are written to long columns, this can cause performance issues both with data writing and the inability to index long columns. In

EDQ 12.2.1.4.3, long text handling has been refactored to support specific designation of 'long' columns in snapshots and staged data, such that all other columns can be written to efficiently and indexed. For more information, see [Working with Long Text](#).

1.2.1.2 Enhanced Redirect URI Selection for OpenID Authentication

The EDQ OpenID integration framework is enhanced to allow configuration of different redirect URIs, based on the incoming host name.

This allows logins either through a load balancer, or directly to a specific host behind a load balancer, which are otherwise protected via an OpenID framework such as Azure AD. For detailed instructions, refer to [Enabling Multiple URI Redirects for OpenID Authentication](#).

1.2.1.3 Case Management REST API Enhancements

The Case Management REST API now allows you to retrieve source and relationship data for alerts and to delete empty cases.

The Case Management REST API includes the following new functionality:

- Include source data in results for alerts: The **getcase** and **runfilter** methods include a new attribute called **sourcedata**. When you retrieve details for alerts, if you set **sourcedata** to true, the results include source and relationship data.
- Delete multiple cases by ID: The **deletecases** method allows you to delete multiple cases/alerts by ID. The user making the call must have the case management bulk delete permission.
- Bulk delete cases using filter: You can filter and delete multiple cases/alerts using the **bulkdelete** method. The user making the call must have the case management bulk delete permission and the permission to use the filter.

For full specification of the API click on [Case Management API Specification](#) in the Web Services menu on the EDQ Launchpad.

1.2.1.4 Support for Amazon Simple Queue Service

EDQ now supports the use of Amazon Simple Queue Service (Amazon SQS) for reading and publishing.

1.2.1.5 Support for User Authentication using Amazon Cognito User Pools

This release includes support for the Amazon Cognito user pool as an identity store.

You can enable SSO for EDQ using OpenID Connect with the OAuth2 authentication method. For detailed instructions, see [Integrating EDQ with Cognito user pools](#).

1.2.1.6 Java Flight Recorder tool to Monitor Web Service Requests

EDQ can generate Java Flight Recording events for web service requests using the standard `jdk.jfr` APIs available in Java 11 and later.

1.2.1.7 Support for Apache Tomcat Version 10

EDQ 12.2.1.4.3 includes support for Apache Tomcat version 10.

Tomcat version 10 and later implement web APIs from Jakarta EE. In Jakarta EE packages are in the `jakarta.*` namespace. The version of the EDQ Tomcat war file built for Jakarta EE is distributed as `jakartaee/edq.war`. Use this file to deploy the EDQ application on Tomcat version 10 or later.

1.2.1.8 Adding Additional JMS Drivers in `<messengerconfig>`

EDQ 12.2.1.4.3 introduces the concept of "named" JMS drivers. The drivers for a particular broker should be installed in `localhome/jmsdrivers/NAME` where NAME is the driver name.

For example, to use a driver for Apache Artemis, create the directory `localhome/jmsdrivers/artemis` and copy `artemis-jms-client-all-2.x.y.jar` to this directory. To refer to a named driver in a JMS messenger configuration, add the line `driver = drivename` to the `<messengerconfig>` section.

The JMS script library also supports this driver property. If the driver property is not specified, the open-source ActiveMQ message broker that is bundled with EDQ is used.

To connect EDQ in Tomcat to JMS in WebLogic, create the directory `localhome/jmsdrivers/wls`, copy the `WLSHOME/wlserver/server/lib/wljmsclient.jar` and `WLSHOME/wlserver/server/lib/wlthint3client.jar` files to this directory, and then refer to the driver in the `<messengerconfig>` section. For example,

```
<messengerconfig>
  java.naming.factory.initial      =
weblogic.jndi.WLInitialContextFactory
  java.naming.security.principal  = weblogic
  java.naming.security.credentials = password
  java.naming.provider.url        = t3://host:8001
  connectionfactory               = jms/factory1
  destination                      = jms/queue1
  driver                          = wls
</messengerconfig>
```

1.2.1.9 Apache Tomcat Clustering Support for EDQ

With EDQ 12.2.1.4.3 you can configure a clustered environment using Apache Tomcat in the same way as EDQ operates when deployed to a WebLogic cluster.

EDQ Clustering on Tomcat uses Coherence Community Edition, which is open source. You can configure clustering on Tomcat using system properties and environment variables. You can use an operational override XML file for more complex cases. For detailed information, see EDQ in a Tomcat Cluster.

1.2.1.10 Connectivity to Apache Parquet Data Stores

Apache Parquet is a data serialization and columnar storage format often associated with Hadoop Ecosystem components such as Hive and Spark. For this release, a new data store is available to enable reading data from a Parquet file. This data store can read multiple files from the landing area and treat them as a single source.

For Parquet columns that contain nested structures, the data store attribute is formed by concatenating the names with colons. For example, `roll_num:min,roll_num:max`,

`roll_num:count` and so on. Note that EDQ does not support columns that can include multiple values, such as nested columns that have repeated values. Such columns will not be displayed in a snapshot definition.

1.2.1.11 Configuring Additional HTTP Response Headers

The OWASP Secure Headers Project (also called OSHP) recommends HTTP response headers that you can use to increase browser security. In EDQ 12.2.1.4.2 and earlier you would have apply a new patch every time a recommended header was missing from EDQ HTTP responses.

EDQ 12.2.1.4.3 onwards you can include any such missing headers by defining them in the *director.properties* file. To configure additional response headers, create a file containing a JSON object in which the attribute names are the header names and the values are the header values. To remove a default header from the response, set the value to **null**.

Add the following to *director.properties*:

```
http.responseheaders = name of the JSON file
```

If the value here is not an absolute file name, EDQ will look for the file in the local configuration directory.

For example, to remove the default **X-Frame-Options** header and add the **X-new-header** header, create a JSON file named *headers.json* that includes the following:

```
{ "X-Frame-Options" : null,  
  "X-new-header"    : "some new value"  
}
```

Now add the following to *director.properties*:

```
http.responseheaders = headers.json
```

To remove all of the default headers, add the following to *director.properties*:

```
http.responseheaders.replace = true
```

1.2.1.12 Tool to Update Database Passwords in *director.properties*

Schema passwords are often stored in an encrypted format in the *director.properties* file, which can be difficult to update when schema passwords are refreshed. EDQ 12.2.1.4.3 includes a new tool called `setpws.jar` that you can use to update EDQ configuration with new passwords for configuration and results schemas.

To see a usage summary, run the following command:

```
$ java -jar setpws.jar -help
```

For EDQ instances that are created from the Oracle Cloud Infrastructure, this tool also refreshes wallets for Autonomous Database schemas. On such instances, run the tool using the following command:

```
$ sudo -u tomcat /opt/java/bin/java -jar /opt/edq/edq/oracle.edq/  
setpws.jar ...
```

For detailed information about the tool, see Updating Database Passwords using setpws.jar.

1.2.1.13 EDQ Application Launchers for Mac Operating Systems

EDQ 12.2.1.4.3 includes application launchers for Intel-based Mac client operating systems. If you do not want to install Java 1.8 on the client machine you can use these launchers in place of Java Web Start to download and start the client applications.

The Mac OS X launcher is bundled with EDQ 12.2.1.4.3 as a .dmg file. To install on Mac OS X, open the .dmg file and drag the launcher to the Applications folder. The Mac launcher is Intel-based, but can also run on Apple M1 and Apple M2 ARM-based systems.

Note that the application launcher is not enabled by default. To enable the launcher for all users, add the following to `director.properties`:

```
launchpad.mode = xlaunch
```

If the client machine has Java 1.8 installed, you can allow users to choose between the legacy Java Web Start and the EDQ application launcher. To do this, add the following to `director.properties`:

```
launchpad.mode = choose
```

To choose between Java Web Start and the application launcher, right-click the application icon in the Launchpad and select **Launch Preferences ...**

1.2.1.14 Support for Downloading Multiple Files from Cloud Storage

EDQ 12.2.1.4.3 includes a new *Download Multiple Files from Cloud Storage* external task type that you can use to download files from OCI Object Storage, Amazon Simple Storage Service (S3), Azure Storage, and Google Cloud Storage.

The Bucket URL defines the location of the storage bucket. You can enter the URL in any of the following formats:

- **OCI Object Storage:** `https://objectstorage.region.oraclecloud.com/n/tenancy/b/bucketname/o`
- **Amazon S3:** `https://bucketname.s3.region.amazonaws.com`
- **Azure Storage:** `https://account.blob.core.windows.net/containername`
- **Google Cloud Storage:** `https://storage.googleapis.com/storage/v1/b/bucketname/o/`

1.2.1.15 Azure AD Authentication Support for Autonomous Database

EDQ 12.2.1.4.3 supports Azure AD authentication for Autonomous Database Data Store connections using Azure OAuth2 access tokens. For more information, see [Use Azure Active Directory \(Azure AD\) with Autonomous Database](#).

1.2.1.16 Utility to Determine the MIME type of Uploaded Files

Apache Tika is now used to determine the MIME type of uploaded files.

To determine the type for a set of input files, EDQ now includes a new utility called *filetyper* that displays the type as determined by Apache Tika. Use the following command to use the utility:

```
$ java -jar dbtools.jar filetyper file1 file2 ....
```

The file extension is not used to determine the type. Instead, each file is copied to a temporary location and examined to determine the file type.

For example,

```
$ java -jar dbtools.jar filetyper agent.properties mcusers.parquet
MyIcon.icns fake.pdf image.jpg server.xml
agent.properties: text/plain
mcusers.parquet: application/x-parquet
MyIcon.icns: image/icns
fake.pdf: application/x-msdownload
image.jpg: image/jpeg
server.xml: application/xml
```

Here, the file *fake.pdf* is a Windows executable file, which is renamed with the .pdf extension. The output shows that the file type extension is not used to determine the file type.

1.2.2 Applying a Bundle Patch for Oracle Enterprise Data Quality

A bundle patch is an official Oracle patch for Oracle Fusion Middleware components on baseline platforms. Each bundle patch includes the libraries and files that have been rebuilt to implement one or more fixes. All of the fixes in the bundle patch have been tested and are certified to work with one another.

Note:

Oracle recommends that you have the latest version of Opatch (version 13.9.2.0.0+ or later) from My Oracle Support. Opatch requires access to a valid Oracle Universal Installer (OUI) Inventory to apply patches.

Oracle Fusion Middleware 12.2.1 products are installed with OPatch NextGen 13.9.2.0.0 to apply interim patches. The Oracle patch mechanism (Opatch) is a Java-based utility that runs on all supported operating systems. Opatch requires installation of the Oracle Universal Installer.

For detailed instructions, refer to [Manually Patching Oracle Enterprise Data Quality](#).

1.2.3 Issues Resolved for 12.2.1.4.3

This section describes issues resolved in this release.

Table 1-2 Issues Resolved

Issue	Notes
OpenID login does not work in a cluster	OpenID logins such as Azure AD, IDCS, and Cognito failed if the server was running in a cluster. This issue has been resolved in this release. The pending login object is not serializable, but is stored in a cluster-wide map.
IllegalMonitorStateException error while polling for messages from an OCI stream, Kafka topic or SQS queue	If an error occurred while polling for messages from an OCI stream, Kafka topic, or an SQS queue, EDQ would throw the following error: <pre>java.lang.IllegalMonitorStateException: current thread is not owner</pre> Additionally, a Kafka receiver could fail with an exception when the process was canceled. The poll call would throw a <code>WakeupException</code> , which was treated as an error. These issues have been resolved in this release.
Running a match process can fail with class not found exceptions	On EDQ servers that were using Java 11 or Java 17, running a match process could fail with class not found exceptions. This happened because the widget manager was optimized to open the jar file only once during the initial processing, which required that the class loader methods were synchronized. For Java version 11 and later a new method was needed, which did not get synchronized. This would cause the jar file to close even if another thread was using it. This issue has been resolved in this release.
Loading list of filters in the Case Management user interface is inefficient	When the Case Management user interface opens, EDQ loads the user and global filter lists from the database. This process was time-consuming because EDQ would decode the entire filter XML document using source name and source data attribute lookup. This issue has been resolved in this release.
Slow Oracle snapshots can fail with timer already cancelled error	To handle data read timeouts the Oracle JDBC driver uses a timer task, which attempts to interrupt the reading thread. In a snapshot, the reading thread is one of the new execution threads that provides better cancellation support. Attempts to interrupt such threads fail with an <code>UnsupportedOperationException</code> error. When a timer task returns this error, the timer itself is cancelled. Any subsequent use of the timer would fail with a <code>Timer already cancelled</code> error. Since all instances of the Oracle driver share a common timer this would cause other snapshots to fail immediately. This issue has been resolved in this release.

Table 1-2 (Cont.) Issues Resolved

Issue	Notes
Charset in Content-Type header is not handled on Tomcat	<p>In EDQ 12.2.1.4.1 web service call requests would be rejected with an HTTP error 415 (invalid media type) error if the Content-Type header included a charset, for example:</p> <pre>Content-Type: text/xml; charset=utf-8</pre> <p>This issue has been resolved in this release.</p>
Internal user lock out fails to work even after multiple invalid login attempts	<p>The security settings for internal users include a mechanism to block a user temporarily or permanently after a configured number of invalid login attempts. This failed to work because the user audit mechanism that would count invalid login attempts failed to include the user ID. This issue has been resolved in this release.</p>
Downloads from Oracle Object Storage or Azure AD would fail	<p>The OCI request signing string requires the date header to conform to the standard locale-specific SimpleDateFormat . In recent versions of Java, the date is formatted differently, which is rejected by OCI. Note that this only happens for dates falling in the month of September. This issue has been resolved in this release.</p>
Cannot start remote flight recordings on systems running Java version 11 and later	<p>On systems that run Java version 11 or 17, EDQ does not expose the JFR MBean (<code>jdk.management.jfr:type=FlightRecorder</code>) in the JMX server. This MBean is used by tools such as Java Mission Control to initiate the flight recording. This issue has been resolved in this release. Now the set of platform MXBeans is determined automatically by using the ManagementFactory methods.</p>
<code>configapp.jar</code> fails due to missing classes	<p>The <code>configapp.jar</code> application that is used for setting up Tomcat installs failed because of missing classes. This happened because the service location mechanism now refers to classes in the expressions package. This issue has been resolved in this release.</p>
Case Management REST API returns invalid JSON in source data	<p>When case data was generated by the Case Management REST API with the source data option, invalid numeric values could appear in the JSON as NaN. For example:</p> <pre>"dnRiskScorePEP": NaN#</pre> <p>This issue has been resolved in this release.</p>
Case Management user interface can crash when running a large number of case sources	<p>The Case Management user interface could fail to respond when running or editing filters with a large number of sources selected, for example more than 200 case sources or possible states. This issue has been resolved in this release.</p>
Database exports may attempt to create tables with duplicate column names	<p>When the "create table" option was selected for a database export, the EDQ framework attempted to create unique column names by truncating the names and adding a numeric suffix. This could result in duplicate column names. This issue has been resolved in this release.</p>

Table 1-2 (Cont.) Issues Resolved

Issue	Notes
Location attribute is "unspecified" in WSDLs generated on Tomcat	<p>In WebLogic and earlier versions of Tomcat, the web service endpoint for a web service request was set to:</p> <pre data-bbox="769 405 1344 464"><soap:address location="http://server/edq/webservices/rt:summer"/></pre> <p>However, a WSDL generated for a web service request on Apache Tomcat contained "unspecified" for the address. For example:</p> <pre data-bbox="769 569 1289 590"><soap:address location="unspecified"/></pre> <p>This issue has been resolved in this release.</p>
No messages logged if an invalid external user realm type is specified	<p>In <i>login.properties</i> if you specified an invalid realm "type", for example the value contained a trailing space, no message were generated in the logs. This made it difficult to diagnose issues. This issue has been resolved in this release.</p>
EDQ application launcher does not support TLS v1.3	<p>The EDQ application launcher failed with a handshake error if the destination server used TLS v1.3 for HTTPS connections. This happened because the launcher did not include the <code>jdk.crypto.ec</code> module that is required for TLS v1.3 support. This issue has been resolved in this release.</p>
Remove unsafe-inline JavaScript content policy	<p>EDQ 12.2.1.4.3 generates an HTTP Content-Security-Policy response header containing:</p> <pre data-bbox="769 1062 1333 1146">default-src 'self' 'unsafe-eval' 'unsafe-inline'; img-src 'self' data:; child-src 'none'; object-src 'none'</pre> <p>While allowing <code>unsafe-inline</code> for scripts is undesirable, removing it caused the swagger REST API pages to not display correctly because these pages contain an inline script that configures the main JSON specification file. Additionally, links in help pages did not work because these pages use <code>href="javascript:golink(...)"</code>, which again counts as inline JavaScript. This issue has been resolved in this release. Inline JavaScript has been removed from the Swagger API docs and the help system.</p>
Restore "login as" icons in EDQ Launchpad with Kerberos	<p>EDQ configurations that use Integrated Windows Authentication (Kerberos) allow users to login to applications automatically. In EDQ 12.2.1.4.0 and earlier versions, the EDQ Launchpad included a "Login as" icon to allow users to manually login to the applications. This was removed in EDQ 12.2.1.4.1. The "Login as" icon is now included in the EDQ Launchpad.</p>
Upgrade rhino 1.7.11 to rhino_1.7.13.0.0.jar	<p>The Rhino JavaScript engine has been upgraded to version 1.7.13.</p>

Table 1-2 (Cont.) Issues Resolved

Issue	Notes
Default read timeouts for Autonomous Database data stores are too low	The data store definition for Oracle Autonomous Database sets the default read and connect timeouts to 20 seconds. This is intended to prevent long delays if a proxy is required but has not been configured. However, when reading from a large external table, the reads may fail with a socket timeout error. This issue has been resolved in this release. You can now configure the default timeouts for each data store.
External lookups can fail if data store specifies a schema	On Snowflake and SQL Server configurations if a schema is set when a data store is defined such that table names are filtered, then external lookups using such data stores may fail. This happens because the SQL generated for the lookup does not include the schema name. To avoid this issue, leave the schema empty when you define the data store.
Oracle data store type does not support BINARY_DOUBLE columns	A BINARY_DOUBLE column in Oracle is reported in JDBC metadata as follows: <code>data type = 101 type name = BINARY_DOUBLE</code> However, the numeric data type value 101 is not a <code>java.sql.Types</code> constant and so the column was not accepted for snapshots. This issue has been resolved in this release.

1.2.4 Known Issues and Workarounds

This section details known issues in this release, and their workarounds.

1.2.4.1 Using Tomcat with Java 8 Causes EDQ to Fail

EDQ 12.2.1.4.3 fails when run on Tomcat (all versions) with Java 8 (JDK 1.8). If you want to run EDQ with Java 8 you need to create an empty `sysdefs/coherence.properties` file in the EDQ local home configuration directory.

Oracle recommends that you run EDQ on Tomcat with Java versions 11 or 17.

1.3 Release 12.2.1.4.2

This document contains release information for Oracle Enterprise Data Quality (EDQ) (12.2.1.4.2) and includes the following sections:

1.3.1 New Features and Improvements

This release introduces the following enhancements:

1.3.1.1 Support for User Authentication using Azure Active Directory

This release includes support for Azure Active Directory (AD) as an identity store. You can enable SSO for EDQ using Azure AD and OpenID Connect.

For detailed instructions, refer to Integrating EDQ with Azure Active Directory.

1.3.1.2 Support for User Authentication using Oracle Identity Cloud Service

You can now integrate EDQ with Oracle Identity Cloud Service (IDCS) for authentication.

For detailed instructions, refer to Integrating EDQ with Oracle Identity Cloud Service.

1.3.1.3 Enhancements to the Apache Hive™ Connector

The Apache Hive connector has been enhanced to support SSL and HTTP transport mode.

For more information, see the Data Stores section of Oracle Enterprise Data Quality Online Help.

1.3.1.4 EDQ Application Launchers for Supported Client Operating Systems

EDQ 12.2.1.4.2 and later releases now include application launchers for supported client operating systems. If you do not want to install Java 1.8 on the client machine you can use these launchers in place of Java Web Start to download and start the client applications.

Click the **Download application launchers** link at the bottom-right of the Launchpad to view the list of available application launchers.

Note that the application launcher is not enabled by default. To enable the launcher for all users, add the following to *director.properties*:

```
launchpad.mode = xlaunch
```

If the client machine has Java 1.8 installed, you can allow users to choose between the legacy Java Web Start and the EDQ application launcher. To do this, add the following to *director.properties*:

```
launchpad.mode = choose
```

To choose between Java Web Start and the application launcher, right-click the application icon in the Launchpad and select **Launch Preferences ...**

1.3.1.5 Connectivity to Snowflake

Snowflake is a data warehousing service that is built on top of the Amazon Web Services or Microsoft Azure cloud infrastructure. EDQ can connect to Snowflake using the standard JDBC driver that Snowflake makes available.

For more information, see the Data Stores section of Oracle Enterprise Data Quality Online Help.

1.3.1.6 New Case Management API

Case Management now includes a REST API to allow filter execution and case updates.

For a specification of the API click on Case Management API Specification in the Web Services menu on the EDQ Launchpad.

1.3.1.7 Support for Java versions 11 and 17

Where EDQ is deployed on Apache Tomcat, Java versions 11 and 17 may be used on the EDQ server.

1.3.1.8 Improved Reporting for Case Management Filter Execution

You no longer need to go through the debug logs to extract information about Case Management filter execution. EDQ now includes trigger points that are called at the start and end of each filter execution. Information about the filter execution is passed to the triggers, which can then be sent to a streaming service such as Java Message Service (JMS), Oracle Cloud Infrastructure Streaming service, and Amazon Kinesis Data Streams.

1.3.2 Applying a Bundle Patch for Oracle Enterprise Data Quality

A bundle patch is an official Oracle patch for Oracle Fusion Middleware components on baseline platforms. Each bundle patch includes the libraries and files that have been rebuilt to implement one or more fixes. All of the fixes in the bundle patch have been tested and are certified to work with one another.



Note:

Oracle recommends that you have the latest version of Opatch (version 13.9.2.0.0+ or later) from My Oracle Support. Opatch requires access to a valid Oracle Universal Installer (OUI) Inventory to apply patches.

Oracle Fusion Middleware 12.2.1 products are installed with OPatch NextGen 13.9.2.0.0 to apply interim patches. The Oracle patch mechanism (Opatch) is a Java-based utility that runs on all supported operating systems. Opatch requires installation of the Oracle Universal Installer.

For detailed instructions, refer to [Manually Patching Oracle Enterprise Data Quality](#).

1.3.3 Issues Resolved for 12.2.1.4.2

This section describes issues resolved in this release.

Table 1-3 Issues Resolved

Issue	Notes
Database errors can cause match process to fail and leave case source locked	Match processes lock the case source during the case generation phase. Any database errors that occurred during match processing could cause the process to fail and lock the source indefinitely. This could then cause subsequent case generation processes to fail. This issue has been resolved in this release. Now any such source locks that happen due to database errors are released during processing.
Case Management filter search returns null pointer exception if no source data found	If you have enabled SQL searches and run a Case Management filter before any source data is created, the query builder generated a null pointer exception. This issue has been resolved in this release. Now the search run completes with no results returned.
Web Service Tester tool results show incorrect results in Google Chrome	If you started the Web Service Tester tool using the Google Chrome browser, the tool displayed empty results incorrectly. This issue has been resolved in this release.
Errors in Groovy trigger scripts cause application failures	The script engine failed to catch exceptions generated by triggers in Groovy scripts, which in turn caused application errors. This issue has been resolved in this release.
EDQ cannot connect to a Teradata database using JDBC driver version 16	EDQ connections to a Teradata database using JDBC driver version 16 failed with an "invalid connection parameter name" error. This issue has been resolved in this release.
EDQ internal user store fails to use pre-defined group mappings	The EDQ internal user store lists all the users in each EDQ group and uses the pre-defined group mappings in <i>login.properties</i> . The user store ignored these pre-defined mappings when determining group membership. This issue has been resolved in this release.
Index creation fails for columns that exceed the maximum index key length of 6398 bytes	The total index length is computed as the sum of the width of all indexed columns plus the number of indexed columns. If you tried to create an index for columns that exceeded the 6398 bytes key length, indexing failed with an "ORA-01450: maximum key length (string) exceeded" warning message. These warning messages could flood the EDQ logs after job completion. This has been fixed. EDQ does not attempt to create an index if the key length exceeds the maximum limit.
User administration page cannot display large number of users	The user administration page displayed the list of users in a scrollable table without paging. However, if the total number of users exceeded 500 the page would fail to display all the users. This issue has been resolved in this release. You can now filter the list to view a subset of users on the page.
Reader processors may not terminate if error encountered while reading Kafka records	If a message decoding error occurred when reading records from a Kafka topic, the process would not terminate automatically. You had to cancel it manually. This issue has been resolved in this release.

Table 1-3 (Cont.) Issues Resolved

Issue	Notes
Case Management reports execution with date ranges can return too many items	If you used SQL mode to run a Case Management report execution using a date value as the row and there were matching items with no value for the date, you would see all dates starting from Jan 1970 and a null value as the second row. Sometimes, these additional rows could cause the report generation to fail due to the large number of items. This issue has been fixed. The date range now starts with the lowest matching date to the specified value.
Scripts that support auto-close mechanism to free resources	If you had to send JMS, OCI, or Kafka messages using a Script Trigger, you would have to establish a JMS connection for each invocation. This is no longer needed. The trigger scripts now create a connection during the initialization phase. These scripts also close the stream and release resources when the trigger is reloaded.
StartTLS encryption for LDAP connections does not verify certificate fully	EDQ applies relaxed checks to the server certificate if you configure encryption for connections between EDQ and LDAP by setting <code>ad.ldap.security = tls</code> in <i>login.properties</i> . For better security, use the following setting instead: <code>ldap.security = tlsverify</code>
Channel binding for LDAP connections	This allows EDQ to properly verify the server certificates. EDQ now provides the option to enable channel binding for LDAP connections over SSL or TLS. To enable channel binding for LDAP connections, add "channelbinding" to the security setting in <i>login.properties</i> : <code>ad.ldap.security = tls,channelbinding</code> or <code>ad.ldap.security = ssl,channelbinding</code> or <code>ad.ldap.security = tlsverify,channelbinding</code> Note that you must be running Java 17 on the EDQ server to use channel binding. This is therefore only supported where EDQ is running on Apache Tomcat.

Table 1-3 (Cont.) Issues Resolved

Issue	Notes
Hints used in Oracle SQL queries should be configurable	<p>Oracle SQL queries previously used only the default <code>FIRST_ROWS</code> hint for interactive queries. The hints for interactive and batch queries are now configurable. Use the <code>oracle.interactive.query.hints</code> and <code>oracle.batch.query.hints</code> properties in <i>director.properties</i> to set hints for interactive and batch queries respectively.</p> <p>For example:</p> <pre>oracle.interactive.query.hints = FIRST_ROWS(1000) PARALLEL(8)</pre> <p>Note the following:</p> <ul style="list-style-type: none"> • Hints are specified without the <code>/*+ ... */</code> wrapper in <i>director.properties</i>. • The default for interactive queries is <code>FIRST_ROWS(100)</code>. • The default for batch queries is no hint. • Hints are exposed as attributes on the Data Access > Repository SQL Creator and Data Access > Results SQL Creator MBeans, which allow you to update the hints on a running system without restarting the server. • Updates to the MBeans are not persisted and must be applied to each server in a cluster.
Case Management searches are slow with <code>ROWNUM</code> in <code>WHERE</code> clause	Previously <code>ROWNUM</code> comparison was used to limit the results returned from a Case Management SQL filler execution. With current Oracle database releases, this is inefficient. Use of <code>ROWNUM</code> has been removed.
Case Management search on extended attributes with Negate option does not return nulls	It was reported that when searching on the extended attributes of alerts with the Negate option selected, the search only returned records with different values; it did not include the nulls. This issue has been fixed in this release. For searches on extended attributes, selecting the Negate option now also returns null values.
Case Management search on extended attributes with Null option does not work	It was reported that when searching on the extended attributes of alerts with the Null option selected, the search only returned null values. This issue has been fixed in this release. For searches on extended attributes, selecting the Null option now returns all rows.
Database commit jobs can run in loop after database error	Occasionally, a WebLogic database connection could be marked as "destroyed" after an RAC failover. If this coincided with a "commit" for a database insert, the "retry" code could get stuck in a loop causing high CPU usage without generating any log output. To cancel such jobs you would have to shut down the application server. This issue has been resolved in this issue

Table 1-3 (Cont.) Issues Resolved

Issue	Notes
Do not enable Java Security Manager	Java Security Manager (-Djava.security.manager) used to be enabled by default for new installations of EDQ on WebLogic servers as it was specified in the setStartupEnv.sh script that sets the Java options for the EDQ server startup group. However, it has been reported that this can cause performance issues for some operations. -Djava.security.manager is no longer specified by default in the setStartupEnv.sh script in EDQ 12.2.1.4.2 and later releases.
Call External Web Service does not populate any data in EDQ	The Call External Web Service processor failed to populate HTTP Response Code, HTTP Response Message, and Web Service Error Message. This issue has been resolved in this release.
Case Management application becomes unstable when processing large number of case sources	On a system that had more than 100 case sources or possible states, the Case Management application became unstable when creating or editing filters when a large number of sources were selected. This issue has been resolved in this release.
dbtools.jar and migration.jar do not include orapki jars required for ADB connections	The dbtools and migration jar files installed with EDQ did not contain the additional driver jar files (oraclepki.jar, osdt.jar, and osdt_core.jar) required for ADB connections. As a result these jar files could not be used for manual connections and schema initialization. This issue is now resolved. The additional driver jar files are now installed along with EDQ.
Retry table creation operation after ORA-00060 error	If EDQ is running heavy batch jobs, the database table creation may fail with an "ORA-00060 deadlock detected" error. This is due to a limitation in Oracle database. This has been fixed. Table creation and drops are retried after an ORA-00060 error. The number of retries is controlled by the <code>oracle.deadlock.feature.retry.count</code> property in <code>director.properties</code> . The default retry count is 5. This means that there are a total of 6 table creation attempts including the 5 retries.
Call External Web Service fails if previous request ended in error	If a Call External Web Service processor request failed, the subsequent request would fail with an array index exception, which would in turn cause the entire process to fail. This issue has been fixed in this release. A processor failure no longer causes the entire process to fail.
In a clustered environment applications can crash if connected immediately after server start	In a clustered environment, if an application such as Case Management connects to a server less than ten minutes after the server has started, the connection could fail with a null pointer error. This issue has been fixed in this release.
Match process does not respond to cancellation requests	In some circumstances, a match process with a real-time input would not respond to cancellation requests. This issue has been fixed in this release.

Table 1-3 (Cont.) Issues Resolved

Issue	Notes
Get result info without full query execution in snapshot definition	When you use custom SQL to define a snapshot, the query is executed to obtain result column information. Some queries on large data sets took a long time, which would then result in a client timeout. This issue has been fixed in this release. The <code>getMetaData</code> method on the <code>PreparedStatement</code> object is invoked first. If this method fails, the query is executed as before.
Use asynchronous processing for promulgation events	Applications (such as Director, Case Management, and so on) use a polling API to receive promulgation events. A request is made to the server, which returns a response every 60 seconds or when an event is available. A separate active thread would run on the server at all times for each application, which could cause significant I/O overhead. This issue has been fixed in this release. To avoid this issue, turn on asynchronous processing by adding the following to <code>director.properties</code> : <pre>promulgation.poll.async = true</pre> This enables the promulgation poll call to use asynchronous processing so that no thread is used until a response is available.
Misleading error shown if invalid new password is entered by user	If a user's password expires they need to set the new password when they launch an EDQ application. If the new password did not meet the required password criteria the "User authentication failed" error was displayed, which was not helpful. The error message has been improved in this release.

1.3.4 Known Issues and Workarounds

This section details known issues in this release, and their workarounds.

1.3.4.1 Using Tomcat with Java 17 Causes Snowflake Snapshots to Fail

The Snowflake driver contains third-party code, which attempts to access certain private Java classes and methods. When used with Java 17, snapshots fail with the following error:

```
JDBC driver internal error: Fail to retrieve row count for
first arrow chunk: sun.misc.Unsafe or
java.nio.DirectByteBuffer.<init>(long, int) not available
```

To avoid this issue, add these runtime options to the `bin/setenv.sh` Tomcat file:

```
JDK_JAVA_OPTIONS="--add-opens    jdk.unsupported/sun.misc=ALL-UNNAMED --
add-opens java.base/java.nio=ALL-UNNAMED"
export JDK_JAVA_OPTIONS
```

1.4 Release 12.2.1.4.1

This document contains release information for Oracle Enterprise Data Quality (EDQ) (12.2.1.4.1) and includes the following sections:

- [New Features and Improvements](#)
- [Applying a Bundle Patch for Oracle Enterprise Data Quality](#)
- [Removal Notice](#)
- [Issues Resolved](#)
- [Known Issues and Workarounds](#)

1.4.1 New Features and Improvements

This release introduces the following enhancements:

1.4.1.1 Stored Credentials

The Stored Credentials feature is a new permission-controlled administrator functionality that allows you to store sets of credentials that EDQ can use to access external systems with high security, such as cloud data storage systems, so that you can use them in file download and upload tasks configured in EDQ, and when calling web services that require the same authentication.

For more information, see the Stored Credentials section of Oracle Enterprise Data Quality Online Help.

1.4.1.2 Extended Limit of Oracle VARCHAR Columns in EDQ Results Schema

Oracle databases 12c and later can be configured to support 32767 bytes instead of 4000 bytes as the maximum size of VARCHAR columns. When EDQ starts up, the maximum size of columns supported by the EDQ repository database is detected automatically and all VARCHAR columns in results tables are created with this size.

For more information, see the Limits in EDQ chapter of Administering Oracle Enterprise Data Quality guide.

1.4.1.3 Data Store for Autonomous Databases

Oracle Autonomous Database is a fully managed, preconfigured database environment with two workload types available, Autonomous Transaction Processing and Autonomous Data Warehouse. For this release, a new data store is available to enable reading data from, and exporting data to, Autonomous databases. This new data store downloads the required wallet file needed to authenticate automatically by calling the ADB REST APIs.

For more information, see the Data Stores section of Oracle Enterprise Data Quality Online Help.

1.4.1.4 Connectivity to Apache Kafka

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.

For more information on connecting EDQ to Apache Kafka, see the Using Apache Kafka with EDQ section of *Integrating Enterprise Data Quality with External Systems*.

1.4.1.5 Connectivity to Apache Hive™ using Kerberos Authentication

Apache Hive™ is a data warehouse software which facilitates reading, writing, and managing large datasets residing in distributed storage using SQL. For connecting EDQ to Apache Hive, you have to configure the server running EDQ to support Kerberos. This requires a valid Kerberos configuration file containing the realm used with Hive.

For more information, see the Data Stores section of Oracle Enterprise Data Quality Online Help.

1.4.1.6 Connectivity to Apache Avro™ Data Stores

Apache Avro™ is a data serialization and storage format often used in conjunction with Big Data systems. For this release, a new data store is available to enable reading data from, and exporting data to an Avro file. The schema in Avro files used with the EDQ data store must define the data as a record type. The record field names map to attributes in EDQ snapshots and exports.

For more information, see the Data Stores section of Oracle Enterprise Data Quality Online Help.

1.4.1.7 Connectivity to AWS Redshift

AWS Redshift is a fast, simple, cost-effective data warehousing service. EDQ can connect to AWS Redshift using the standard JDBC driver that AWS makes available.

For more information, see the Data Stores section of Oracle Enterprise Data Quality Online Help.

1.4.1.8 Connectivity to JSON Data Stores

EDQ can read and write data in JSON and JSON Lines formats. The data is represented as an array of objects, either at the top level of the text file, or at a simple attribute path within the file.

For more information, see the Data Stores section of Oracle Enterprise Data Quality Online Help.

1.4.1.9 Connectivity to Oracle Service Cloud

Oracle Service Cloud is a CRM solution that improves customer service and daily operations through the use of its service request management, knowledge base, and customer portal. A new data store is available to enable reading data from Oracle

Service Cloud. The Oracle Service Cloud data store uses a driver that is suitable only for data extraction, so it is not possible to export data for this type of data store.

For more information, see the Data Stores section of Oracle Enterprise Data Quality Online Help.

1.4.1.10 Support for Case Management Filter and Report Execution by using SQL and Oracle Text

This release includes optional support for Case Management filter and report execution using pure SQL, in conjunction with Oracle Text and JSON indices. Through this option, you can enable SQL filtering and disable Lucene. The EDQ repository database must be running Oracle database 12.2 or later to use this option.

For detailed instructions, refer to [Switching to SQL and Oracle Text for Case Management Filters and Reports](#).

1.4.1.11 Support for Reconnect Retries in JMS Connection

This release includes support for allowing retries when connecting to JMS message queues or topics, for either reading (consuming messages) or writing (providing messages). For allowing connection retries, configure the new properties in `messengerconfig` section of the `realtime bucket` XML file.

New properties are:

```
initialattempts: Number of times to try making initial connection. Default = 1
retryattempts: Number of times to try to reconnect after error. Default = 0
retrydelay: Time to wait between connection attempts. Default = 5s
```

The property `retrydelay` can be specified as:

- Nms - N milliseconds
- Ns or N - N seconds
- Nm - N minutes
- Nd - N days

for example, the following settings set the delay to 5 seconds (5000 milliseconds):

```
retrydelay = 5000ms
retrydelay = 5
retrydelay = 5s
```

For more details on JMS, refer to [Using JMS with EDQ](#) documentation.

1.4.2 Applying a Bundle Patch for Oracle Enterprise Data Quality

A bundle patch is an official Oracle patch for Oracle Fusion Middleware components on baseline platforms. Each bundle patch includes the libraries and files that have been rebuilt to

implement one or more fixes. All of the fixes in the bundle patch have been tested and are certified to work with one another.

 **Note:**

Oracle recommends that you have the latest version of Opatch (version 13.9.2.0.0+ or later) from My Oracle Support. Opatch requires access to a valid Oracle Universal Installer (OUI) Inventory to apply patches.

Oracle Fusion Middleware 12.2.1 products are installed with OPatch NextGen 13.9.2.0.0 to apply interim patches. The Oracle patch mechanism (Opatch) is a Java-based utility that runs on all supported operating systems. Opatch requires installation of the Oracle Universal Installer.

For detailed instructions, refer to [Manually Patching Oracle Enterprise Data Quality](#).

1.4.3 Removal Notice

The following section describes the features that have been removed in this release:

- **Removal of Active MQ server:** From this release, all the Active MQ server jars are removed and it is therefore no longer possible to launch or use an embedded Active MQ server. The recommended architecture for Active MQ is to run a separate broker outside EDQ.
- **Removal of ability to view client session logs:** The ability to view client session logs for user sessions has been removed from the Administration pages on the Launchpad.
- **Removal of XML support in EDQ configuration REST APIs:** Currently, XML data can be returned from configuration REST API calls by providing the 'accept: text/xml' header. However, support for XML requires `javax.xml.bind.annotation` classes and this package is not available in Java 11. XML support has therefore been removed from EDQ configuration REST APIs.
- **Removal of `wSDLizer` and support for 'Global' web services:** The `wSDLizer` is used in conjunction with Java XML binding (JAXB) and it has no impact on standard EDQ web services. As part of the transition towards stable Java 11 support, the `wSDLizer` and JAXB processing have been removed from EDQ web services.
- **Removal of FTP server:** The embedded FTP server has been removed from this release due to security reasons. The SFTP interface should be used instead.

1.4.4 Issues Resolved

This section describes issues resolved in this release.

Table 1-4 Issues Resolved

Issue	Notes
Error when saving case attachments containing multibyte characters	An issue was reported in which the file name attached to a case is truncated to 80 characters, when the first 80 characters contained multibyte characters. This issue has been fixed and it is now possible to attach file names containing multibyte characters without truncation. Also, it is to be noted that only PDF attachments are supported due to security reasons.
Error using decision input on match processor	An issue was reported in which an error occurred when using the decision input on the match processor to update alerts. It was noted that the process failed when the decision phase components were enabled. This issue has been fixed in this release.
Reference data is not updated correctly when written to and then read in the same job	An issue was reported in which the reference data in a reader is not refreshed, when the reference data is changed in the same running job. This issue has been resolved in this release.
Dashboard Administration may not work through a load balancer to a cluster	When using a load balancer in a cluster containing two or more managed servers, it was reported that Dashboard Administration fails to start with an 'Invalid Credentials' error message. This issue has been fixed and Dashboard Administration now starts successfully from behind a load balancer.
Results for a job and run label are not displayed in the Server Console	An issue was reported in which the results for a job and run label are not displayed in the Server Console. For example, the same job in the same project which was run with run label 'Test1' and then 'Test2' shows both sets of results initially, but only the results for run label 'Test2' after the same user logs out and then back in to the Server Console. Whenever the user logs in again, only the latest run label for the job is displayed and there is no way to display the others. This issue has been resolved in this release.
Spurious "Access to PID X is strictly forbidden" errors seen in Server Console and logs	It was reported that on logging in to the Server Console as a non-administrative user and running a job with a run label in a project for which the non-administrator user does not have access (that is, it is marked as 'Administrators' only), an error message "Access to PID X is strictly forbidden" is displayed.
If a "run everywhere" job is cancelled, it still runs when other servers start	An issue was reported in which a cancelled "run everywhere" job remains in the list of running jobs, ready to run when other servers are started. This issue has been resolved in this release.
Threads used for "run everywhere" jobs on a new server do not have correct context	When a server in a cluster starts up, it will automatically run any "run everywhere" jobs which are already running in the cluster. The threads used for these jobs do not have the correct context class loader set and processes may fail to locate resources loaded through <code>Thread.currentThread().getContextClassLoader()</code> . This can cause problems with the Groovy <code>JsonSlurper</code> since this requires a fast string service loaded using the context class loader. This issue has been resolved in this release.

Table 1-4 (Cont.) Issues Resolved

Issue	Notes
JSON decoding fails with JMS messages	An issue was reported in which JSON decoding fails and displays an error when receiving JMS messages. This issue has been resolved in this release.
Case state expiry can severely degrade indexing speed	An issue was reported in which case state expiry severely degrades the indexing speed. For example, consider a workflow with a state which has an expiry time set to 1 minute. If a large number of new cases/alerts are generated or updated to this state, a large number of expiries will happen and the whole indexing process will essentially stall. This issue has been resolved in this release.
Database connection framework does not support filtering by catalog	An issue was reported in which filtering by catalog is not supported in EDQ. Snowflake maps database names to catalog names. So for a Snowflake connection, tables from all databases are listed. Attempting to read from a table which is not in the selected database results in an error. This issue has been resolved in this release.
Unnecessary synchronization in Runtime progress manager	An issue has been reported in which there is unnecessary synchronization in the runtime progress manager. This can cause a dramatic slowdown in performance in a cluster. This issue has been resolved in this release.
Case creation/update places excessive load on lock manager	Cases and alerts are created/updated in batches of 30. For each batch, a lock request containing 60 GUIDs is made. This can cause significant performance problems in a clustered environment where locks are stored in Coherence. This issue has been resolved in this release.
Allow selective disabling of case management indices	<p>An issue was reported in which there is a need to disable certain case management indices in order to improve performance. There are mainly five case management indices as follows:</p> <ul style="list-style-type: none"> • Case: This is essential and takes the least time to build. • Supplementary Data: This is important to enable detailed searching, and takes some time to build. • Comments: This is useful for searching on previous comments, and takes some time to build. • Transitions: This is not needed for many use cases, and is usually cheap to build. • History: This is not needed for many use cases, and is very expensive to build. <p>To disable one or more indices, you can add the <code>cm.disabled</code> property to <code>director.properties</code>. The value should be a comma/space separated list of index names. The index names are <code>history</code>, <code>sd</code>, <code>comment</code>, and <code>transitions</code>. You can also add the value <code>"reporting"</code> to hide the reporting piece in the case management user interface. This does not affect indexing. The case index cannot be disabled.</p> <p>For example, to disable the history index, you can set the property as <code>cm.disabled = history</code>.</p>

Table 1-4 (Cont.) Issues Resolved

Issue	Notes
Names of temporary tables created during case management may clash in a cluster	Case management creates temporary tables to collate supplementary data and other data. These tables are named with a prefix and a per-server atomic integer. In a cluster, this may lead to clashes in table names with subsequent chaos. This issue has been resolved in this release.
JMS reception does not fail even if message broker stops	An issue has been reported in which if, a process is reading from a JMS queue or topic, and the message broker fails, the process does not error. This issue has been resolved in this release, so that real-time processes reading from JMS queues or topics will fail if the broker becomes unavailable, unless connection retries have been configured and are successful.

1.4.5 Known Issues and Workarounds

This section details known issues in this release, and their workarounds.

1.4.5.1 Opening Server Package File in Apple OSX Causes Error

An issue has been reported in which a user right-clicked on the Server node in the Project Browser (in Director) to open a Server Package file, and received a null pointer exception. In the meantime, opening Server Package file can still be accomplished in the following ways:

- **File > Open Package File.** Then navigate to the Dxi file containing project elements and import the file. This can then be copied into the Project Browser.
- Or
- Drag and drop.

1.5 Release 12.2.1.4.0

This document contains release information for Oracle Enterprise Data Quality (EDQ) (12.2.1.4.0) and includes the following sections:

- [New Features and Improvements](#)
- [Issues Resolved](#)
- [Known Issues and Workarounds](#)
- [Upgrade Considerations](#)

1.5.1 New Features and Improvements

This release introduces the following enhancements:

1.5.1.1 Call External Web Service Processor

For this release EDQ offers a new processor called **Call External Web Service** which takes input data, configuration information, request payload and parses web service responses as

the processor output. This processor makes it easier to call external REST web services within EDQ processes.

For more information, see the Call External Web Service section of *Oracle Enterprise Data Quality Online Help* .

1.5.1.2 Cassandra Data Store

A new data store is available to enable reading data from, and exporting data to, Cassandra databases.

1.5.1.3 Multiple Value Reference Data Editing

This release enables multiple value editing of Reference Data entries, to improve productivity when editing reference data, especially for standardization and product data classification purposes.

1.5.1.4 New REST API Documentation and Testing capabilities

New documentation of EDQ's REST APIs is available from the EDQ Launchpad. The new documentation gives details and examples for all services, and a built-in testing facility. The 'fixed' services to manage EDQ configuration and jobs are documented under Web Services – REST API Specification, and the dynamic services, built by configuration, are documented under Web Services – Web Services - Web Services REST Endpoints.

1.5.2 Issues Resolved

This section describes issues resolved in this release.

Table 1-5 Issues Resolved

Issue	Notes
Case management reports with aggregations are not consistent with the drilldowns	Issues of consistency and overlap in some reports were noted. These issues have been fixed in this release.
Event log 'EventId' column removal can cause EDQ 9.0 upgrade issue	Users of version 9 of EDQ may have selected all event log columns (including <i>EventId</i>) for display and saved this selection. EventId is no longer used, and users of previous releases reported that trying to load an event log that contains the <i>extra</i> saved column caused the application to shut down. That condition has been fixed in this release.
Client/Server Timezone issues when using filters in Case Management	Conflicting time stamps (server versus client) have been resolved.
Unable to open 'Extract Building Identifier' parser in CDS Standardize Address	Loading symbol data was causing Director to run slowly in some circumstances. The issue has been resolved.
Vertical Data view in Case Management not appearing	An issue preventing proper display of Alerts <i>in the vertical view</i> was fixed in this release.

Table 1-5 (Cont.) Issues Resolved

Issue	Notes
Incorrect key generation export name in CDS run profile	An incorrect key generation export name was discovered and fixed in this release.
IndexOutOfBoundsException error on job email icon with duplicate user display names	Duplicate display names were causing an error to be displayed when users clicked the job email icon. The issue has been resolved in this release, and job notification emails are generated without error.
WebService publishing fails if HTTPS is enabled but HTTP is disabled	An error occurred when EDQ's managed server was configured so that the main listen port was disabled but the SSL listen port was enabled. This issue has been resolved in this release.
Re-enable On-line Help for non-English languages	In a previous release, the online Help was available only in the English language. It is available in 10 languages (including English) in this release.
Match options for data-only from contributing comparisons are not saved	It was discovered that the match options for <i>Data from only contributing comparisons</i> and <i>Data from only contributing compound comparisons</i> were not being saved. It has been corrected in this release.
Download task does not support HTTPS through a proxy server	Using an HTTPS URL in the download task was causing an error when the task passed through a proxy. This issue has been resolved in this release.
Flag Key changes not reflected during Case Source import	An issue was reported in which Flag Key changes were not visible after import. The issue was resolved in this release.
Match cluster limit warnings rendered incorrectly in HTML job log report	An issue in which Match cluster limit warnings were being incorrectly rendered in the Event Log Job Messages HTML output has been corrected in this release.
Correct error status not returned when running runjob	Error status was not correctly reported (no error returned) when running runjob. This issue has been fixed and errors are now reported correctly.
E-mail notification does not support SMTP authentication	An issue was reported in which the email notification function did not work when an SMTP server requiring authentication was used. The defect has been addressed, and notifications are working properly in this release.
File uploads do not work with built-in SFTP server	When connecting to the internal SFTP server, file uploads were returning an error (<i>Couldn't get handle</i>). The issue has been fixed in this release.
Length of keys generated in batch can exceed staging table column size	An issue was reported (in Customer Data Services Pack) in which errors occurred when the length of keys generated by EDQ exceeded the width of the table column. The keys are now truncated to address this issue.
Case Management Reports use all values that exist, rather than those in filter	Case Management Reports were erroneously including placeholders (rows/columns) for values not called for in the filter, in addition to the data called for in the filter. This issue has been fixed in this release.

Table 1-5 (Cont.) Issues Resolved

Issue	Notes
Case Management temp files not deleted until server restart	Temp files were previously being retained until the server was restarted, at which time the temp files were deleted. The application has been updated so that the temp files are cleared when the application is closed.
Issue email notifications do not work if an LDAP userdisplayname is configured	The presence of <i>userdisplayname</i> was interfering with issue email notification. The issue has been resolved, and the presence of <i>userdisplayname</i> does not cause any problems.
Job email notifications do not work if an LDAP userdisplayname is configured	The presence of <i>userdisplayname</i> was interfering with job email notification. The issue has been resolved, and the presence of <i>userdisplayname</i> does not cause any problems.
Excel exports from Case Management reports with aggregations contain <code><nobr></code> tags	<code><nobr></code> and <code></nobr></code> tags were appearing in some Case Management reports exported to Microsoft Excel. This issue has been resolved.
One-way web service returns 'unresolvable error' with webservice tester	When testing a one-way web service using the webservice tester, an error message appeared (<i>unresolvable error</i>). This release includes the fix to this issue.
Configuration Analysis generates a null pointer exception when comparing jobs	An issue was reported in which users received null pointer exceptions when using Configuration Analysis when comparing jobs. This release addresses the issue and it is no longer observed.
Case Management 'State Changed By' filter on user Display Name does not work	When creating a filter using <i>State Changed By</i> , users can now filter users based on their Display Name.
Lucene index update is not committed for user updates of individual cases	Users reported not being able to search on attributes updated in Cases or Alerts until a re-index was performed. Functionality was updated to enable searching on updated attributes without having to re-index.
Option to log Case Management report requests for debug purposes	<p>To provide additional debug logging of all user Case Management report requests:</p> <ol style="list-style-type: none"> 1. Access the following new EDQ mbean in JConsole on a running EDQ server: edq Logging Case Management Filter Execution 2. Invoke the <code>setLevel</code> operation with a value of <i>FINE</i>. Whilst not recommended, this setting can be made permanent by adding the following line to <code>logging.properties</code> and restarting EDQ: <code>com.datanomic.director.casemanagement.se.arch.level = fine</code>

Table 1-5 (Cont.) Issues Resolved

Issue	Notes
Prevent running of very large Case Management reports to avoid memory problems	By default, attempts to run Case Management reports in which either axis contains more than 1000 entries will be rejected, and an error message displayed. The size of this limit is configurable, and can be set using the following new parameter in <code>override.properties</code> , for example to change it to <i>5000</i> : <code>casemanager.maxReportWidth = 5000</code>
Some CM reports with a date field with a week granularity will never finish	An issue has been reported in which running CM reports with a granularity of <i>week</i> caused looping that would not allow the report to be completed. This issue has been resolved in this release.

1.5.3 Known Issues and Workarounds

This section details known issues in this release, and their workarounds.

1.5.3.1 Opening Server Package File in Apple OSX Causes Error

An issue has been reported in which a user right-clicked on the Server node in the Project Browser (in Director) to open a Server Package file, and received a null pointer exception. In the meantime, opening Server Package file can still be accomplished in the following ways:

- **File > Open Package File.** Then navigate to the Dxi file containing project elements and import the file. This can then be copied into the Project Browser.
Or
- Drag and drop.

1.5.4 Upgrade Considerations

This section details major considerations for upgrading to 12c (12.2.1.4.0).

1.5.4.1 Upgrades in an Oracle WebLogic Server Environment

- Read for guidelines for preparing to upgrade to Oracle Fusion Middleware 12c (12.2.1.4.0). This documentation also includes descriptions of terminology changes that you must understand to move forward to a 12c environment.
- If your current EDQ version is 11.1.1.7 or later *and* was installed with Oracle Universal Installer (OUI), you can use the Upgrade Assistant to upgrade your installation directly to 12c (12.2.1.4.0).
- All EDQ components must be shut down and remain stopped until you are prompted to start them at the end of these upgrade instructions. The WebLogic Server console must remain running until you are prompted to shut it down during the upgrade procedure.

For all upgrade scenarios, see *Upgrading Enterprise Data Quality in Installing and Configuring Oracle Enterprise Data Quality* guide.

1.5.4.2 Upgrades in an Apache Tomcat Environment

You can perform a direct upgrade to version 12c of EDQ only from Tomcat version 8. If you are running an earlier version of Tomcat, you must upgrade Tomcat to version 8 before proceeding with the EDQ upgrade. See the Apache Tomcat documentation at

<http://tomcat.apache.org>

To upgrade to 12c (12.2.1.4.0), see Upgrading Enterprise Data Quality in *Installing and Configuring Oracle Enterprise Data Quality* guide.

A

Switching to SQL and Oracle Text for Case Management Filters and Reports

This appendix describes how to use SQL and Oracle Text for Case Management filter and report execution.



Note:

This feature is applicable only for EDQ 12.2.1.4.1 release.

From this release, you can use pure SQL in conjunction with Oracle Text and JSON indices for Case Management filter and report execution. Note that switching from Lucene to Oracle Text is a considerable change and may require tuning to improve search performance. Oracle recommends that you test all aspects of search behavior and performance before you configure your production systems to use Oracle Text.

Follow the below procedure to enable SQL filtering and disable Lucene.



Note:

To use this option, the EDQ repository database must be running Oracle database 19.10 or later.

It includes the following sections:

A.1 Working of Case Management Filters

Case management filters are mapped to SQL searches on the `dn_case` and supporting tables. Filters on the case key, case description and comment text are mapped to Oracle Text searches. Filters on source data are mapped to Oracle JSON Text searches on a new column in the supplementary data table which contains JSON encoded source data. All other fields do not support free text searches and are mapped to simple SQL predicates.

A.2 Understanding Oracle Text Expressions

There are very few differences between Lucene filter expressions and Oracle Text Context Grammar. For more guidance on the Oracle Text search syntax, refer to [The CONTEXT Grammar](#) documentation.

The existing filters use Lucene search syntax and you have to modify them to work correctly with Oracle Text. Asterisk ("*" - wild card) characters in filters are replaced with % automatically.

Refer to [Key differences in Search Functionality between Lucene and Oracle Text in EDQ Case Management](#) for more information.

A.3 Using Oracle Text Options

Indexes created for searches using Oracle Text have a large number of configuration options. The best options to use for any installation greatly depends on the data and typical search patterns.

For example, an index can be enhanced to improve the performance of prefix searches such as A%. The length of the prefix can also be configured. Improving prefix search performance entails an additional cost in index maintenance and storage size. These options can be set on the EDQ repository database housing the EDQCONFIG schema.

EDQ provides a script that sets these options for typical use cases, for example - to include a prefix optimization for prefixes up to 3 characters long when searching on source attributes, but this may be tuned for individual requirements. For more information on this, refer to [Oracle Text Indexing Elements](#).

The performance of comment searches depends on the volume of comment data. You can do one of the following to optimize search performance:

- Uncheck the **Comment** option when using Quick Search.
- Add the following indexes:

```
create index idx_comment_cid on dn_casecomment(case_id)
create index idx_comment_del on dn_casecomment(deleted_flag)
```

A.4 Updating Schema

To enable Oracle Text functionality, you must grant the CTXAPP role for the EDQ configuration schema user.

You must also grant the CREATE JOB system privilege to the user. Use the following SQL commands:

```
GRANT "CTXAPP" TO "USERNAME";
GRANT CREATE JOB TO "USERNAME";
```

where USERNAME is the EDQ configuration schema user name.

A.5 Creating Helper Index and Column

Case and alert permissions are checked directly in the SQL filters. You have to create a new index on the case table to improve performance for these checks:

```
CREATE INDEX idx_dn_case_permission ON dn_case(permission);
```

Source attribute searches are performed using a new JSON data column in the supplementary data table.

Use the following command to add the new column:

```
ALTER TABLE dn_supplementarydata ADD json BLOB CONSTRAINT jcheck CHECK (json IS JSON);
```

A.6 Creating Text Indices

The following SQL script creates the required Oracle Text and JSON indexes:

```
BEGIN
  CTX_DDL.create_preference('dn_textpref', 'BASIC_LEXER');
  CTX_DDL.create_stoplist('dn_textstop', 'BASIC_STOPLIST');
  CTX_DDL.create_preference('dn_wordlist', 'BASIC_WORDLIST');
  CTX_DDL.set_attribute('dn_wordlist', 'PREFIX_INDEX', 'TRUE');
  CTX_DDL.set_attribute('dn_wordlist', 'PREFIX_MAX_LENGTH', '3');
END;
/

CREATE INDEX dn_case_key_text ON dn_case (key_label)
  INDEXTYPE IS CTXSYS.CONTEXT
  PARAMETERS('sync (every "freq=secondly;interval=20") lexer dn_textpref
stoplist dn_textstop wordlist dn_wordlist')
/

CREATE INDEX dn_case_desc_text ON dn_case (description)
  INDEXTYPE IS CTXSYS.CONTEXT
  PARAMETERS('sync (every "freq=secondly;interval=20") lexer dn_textpref
stoplist dn_textstop wordlist dn_wordlist')
/

CREATE INDEX dn_casecomment_text ON dn_casecomment (case_comment)
  INDEXTYPE IS CTXSYS.CONTEXT
  PARAMETERS('sync (every "freq=secondly;interval=20") lexer dn_textpref
stoplist dn_textstop wordlist dn_wordlist')
/

CREATE SEARCH INDEX dn_supp_json ON dn_supplementarydata (json) FOR JSON
PARAMETERS('sync (every "freq=secondly;interval=20") wordlist dn_wordlist')
/
```

The indexes are updated asynchronously every 20 seconds. By default, prefix searches are optimized for prefix lengths up to 3 characters. No special language analysis is enabled.

You can find the script in the `cmsql` directory of the EDQ Home directory after updating to this release.

A.7 Populating JSON

Use the tool `sdjson.jar` to populate the supplementary data JSON column. Execute the following command to run the tool:

```
$ java -jar sdjson.jar oracle:#service@HOST:PORT/USER/PW
```

Tests have shown that the overall time for the conversion and indexing is much less if the JSON index is created before the population step.

A.8 Enabling SQL and Oracle Text Usage for Filtering and Reports

A new setting `cm.filter.sql` in `director.properties` file controls the usage of SQL or Lucene for filters and reports.

```
cm.filter.sql = off
```

Lucene is used for all filters and reports. This is the default for compatibility with earlier versions.

```
cm.filter.sql = on
```

SQL and Oracle Text is used for all filters and reports. Lucene indexing is disabled.

```
cm.filter.sql = optional
```

A Use SQL option is available in the Case Management UI. Lucene indexing is enabled. You can use this setting to compare results and timing between SQL and Lucene searches.



Note:

Use this option only for testing.

A.9 Additional Considerations

The extended attribute (custom flags) columns in the case table do not have database indexes by default. If searches on extended attributes are common without other search filters, it is necessary to create additional indices. This depends on the individual requirements.

B

Manually Patching Oracle Enterprise Data Quality

This appendix describes how to prepare and install the Bundle Patch files for Oracle Enterprise Data Quality.



Note:

This feature is applicable only for EDQ 12.2.1.4.1 release.

You can manually install a patch for Oracle Enterprise Data Quality through OPatch NextGen 13.9.2.0.0 utility.



Note:

Oracle recommends that you always install the latest Bundle Patch.

Patching process uses both unzip and Opatch executables. After sourcing the `ORACLE_HOME` environment, Oracle recommends that you confirm that both of these exist before patching. Opatch is accessible at: `$ORACLE_HOME/OPatch/opatch`

You can check your version using the following command:

```
ORACLE_HOME/OPatch/opatch version
```

Set the `ORACLE_HOME` environment variable to the directory where you have installed Oracle Enterprise Data Quality.

Follow the below procedure to manually install the bundle patch for Oracle Enterprise Data Quality.

It includes the following sections:

B.1 Installing a Patch on Weblogic

Follow the below procedure to install a patch for Oracle Enterprise Data Quality on Weblogic:

1. Unzip the patch zip file to a desired location (for example, `PATCH_TOP`).

```
$ unzip -d PATCH_TOP p28526695_122120_Generic.zip
```

 **Note:**

On WINDOWS, the preferred location is the drive root directory. For example, C:\PATCH_TOP and avoid choosing locations like, C:\Documents and Settings\username\PATCH_TOP. This is necessary due to the 256 characters limitation on windows platform.

On WINDOWS, the unzip command has a limitation of 256 characters in the path name. If you encounter this, use an alternate ZIP utility like 7-Zip to unzip the patch. For example: To unzip using 7-zip, run the command:

```
"c:\Program Files\7-Zip\7z.exe" x p28526695_122120_Generic.zip
```

2. Set your current directory to the directory where the patch is located.

```
$ cd PATCH_TOP/28526695
```

3. Run OPatch to apply the patch.

```
$ opatch apply
```

 **Note:**

When OPatch starts, it validates the patch and makes sure that there are no conflicts with the software already installed in ORACLE_HOME.

In case of opatch conflict, you will see a warning message similar to the one mentioned below:

```
Interim Patch XXXX has Conflict with patch(es) [ YYYY ] in OH ...
Conflict patches: YYYY
Patch(es) YYYY conflict with the patch currently being installed
(XXXX).
If you continue, patch(es) YYYY will be rolled back and the new
patch (XXXX) will be installed.
If a merge of the new patch (XXXX) and the conflicting patch(es)
( YYYY) is required,contact Oracle Support Services and request a
Merged patch.
Do you want to proceed? [y|n]
n
```

You must stop the patch installation and contact oracle support to know how to proceed further.

4. Restart all servers (AdminServer and all Managed server(s)) and ensure that edq.war is successfully deployed. This is necessary to redeploy the original applications and bring the environment back to it's original state.

B.2 Installing a Patch on Tomcat

Follow the below procedure to install a patch for Oracle Enterprise Data Quality on Tomcat:

1. Unzip the patch zip file to a desired location (for example, PATCH_TOP).

```
$ unzip -d PATCH_TOP p28526695_122120_Generic.zip
```

 **Note:**

On WINDOWS, the preferred location is the drive root directory. For example, C:\PATCH_TOP and avoid choosing locations like, C:\Documents and Settings\username\PATCH_TOP. This is necessary due to the 256 characters limitation on windows platform.

On WINDOWS, the unzip command has a limitation of 256 characters in the path name. If you encounter this, use an alternate ZIP utility like 7-Zip to unzip the patch. For example: To unzip using 7-zip, run the command:

```
"c:\Program Files\7-Zip\7z.exe" x p28526695_122120_Generic.zip
```

2. Set your current directory to the directory where the patch is located.

```
$ cd PATCH_TOP/28526695
```

3. Run OPatch to apply the patch.

```
$ opatch apply
```

 **Note:**

When OPatch starts, it validates the patch and makes sure that there are no conflicts with the software already installed in ORACLE_HOME.

In case of opatch conflict, you will see a warning message similar to the one mentioned below:

```
Interim Patch XXXX has Conflict with patch(es) [ YYYY ] in OH ...
Conflict patches: YYYY
Patch(es) YYYY conflict with the patch currently being installed (XXXX).
If you continue, patch(es) YYYY will be rolled back and the new patch
(XXXX) will be installed.
If a merge of the new patch (XXXX) and the conflicting patch(es) ( YYYY)
is required,contact Oracle Support Services and request a Merged patch.
Do you want to proceed? [y|n]
n
```

You must stop the patch installation and contact oracle support to know how to proceed further.

4. Edit `catalina.properties` file, to remove the reference to Oracle Application Development Framework (ADF). Optionally, you can remove `lib.adf` directory from their respective shared library path.

 **Note:**

This step is a one time activity and is specific to 12.2.1.4.1 bundle patch installation.

Remove the entry `${catalina.home}/lib.adf/*.jar` in the file.

For example,

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

5. Stop the Application Server.
6. Deploy the `edq.war` file on your application server.
For more details, refer to the [Tomcat Web Application Deployment](#) documentation.
7. Restart your application server and ensure that `edq.war` is successfully deployed. This is necessary to redeploy the original applications and bring the environment back to it's original state.

C

Reporting for Case Management Filter Execution

This appendix describes how to extract information about Case Management filter execution in Oracle Enterprise Data Quality.

EDQ includes trigger points that are called at the start and end of each filter execution. Information about the filter execution is passed to the triggers, which can then be sent to a streaming service such as Java Message Service (JMS), Oracle Cloud Infrastructure Streaming service, and Amazon Kinesis Data Streams.

The trigger paths are as follows:

- /casemanagement/filter/start
- /casemanagement/filter/end

Run functions in triggers accepting these paths are defined as shown below:

```
function run(path, id, env, json) {  
  ...  
}
```

The **json** argument is a stringified JSON object containing the following attributes:

Attribute	Description
id	Internal filter execution ID
filter	String representation of the filter
type	Possible values are <code>search_filter</code> , <code>report</code> , <code>bulk_update</code> , <code>bulk_delete</code> and <code>export</code>
xaxis	String representation of the report X axis (present for reports only)
yaxis	String representation of the report Y axis (present for reports only)
server	Name of the server where the filter was executed
userid	Internal numeric ID of the user running the filter
user	Name of the user running the filter
userdisplay	Display name of the user running the filter
start	Timestamp of filter start
duration	Filter execution time in milliseconds (present for end calls only)
status	Indicates the reason if the filter did not complete successfully
sql	SQL statement used for the filter (SQL execution mode only)

Attribute	Description
args	Array of SQL bind arguments

The **args** attribute in a report for a SQL execution contains the bind values which replace the '?' placeholders in the SQL text. Each entry in the array contains the following:

type	string, number or date
value	argument value

Example

Suppose a trigger sends filter execution end reports to JMS, an OCI notification service topic, an OCI stream and a Kafka topic.

```

addLibrary("oci");
addLibrary("jms");
addLibrary("kafka");

var oci    = OCI.create("OCI 1")
var topic  = oci.topic("ocidl.onstopic.oc1.phx.aaaaaaa...")
var stream = oci.stream("ocidl.stream.oc1.phx.aaaaaaa...")

var props = config.loadTriggerProperties("jms", "jms\\.properties");
var jms    = JMS.open(Object.assign({}, props,
config.loadCredentials("", props)))
var kprops = config.loadTriggerProperties("kafka", "kafka\
\\.properties");
var kprod  = KAFKA.producer(kprops.topic, kprops)

function getPath() {
    return "/casemanagement/filter/(start|end)";
}

function run(path, id, env, json) {
    if (path.endsWith("end")) {
        topic.publish("filter done", json)
        stream.publish(null, json)
        jms.send(json)
        kprod.publish(null, json)
    }
}

```

Note that the script trigger framework uses one instance for general queries and additional instances for execution. The first instance will never be triggered and thus resources such as JMS connections are not needed. The framework creates a built in

variable `runnable` which is false for the single configuration instance. To optimise resources, this can be used as shown below:

```
var jms = runnable && JMS.open(Object.assign({}, props,
config.loadCredentials("", props)))
...
function run(path, id, env, json) {
  ...
  jms && jms.send(json)
}
```