

Oracle® GoldenGate Big Data and Application Adapters

Release Notes

12c (12.1.2.1)

E61861-11

September 2016

This document contains the release notes for the 12c (12.1.2.1) release of Oracle GoldenGate Application Adapters.

This document contains the following sections:

- [Section 1, "New Features"](#)
- [Section 2, "New or Changed Properties"](#)
- [Section 3, "Deprecated Items"](#)
- [Section 4, "Installation and Upgrades"](#)
- [Section 5, "Corrected Problems"](#)
- [Section 6, "Known Issues"](#)
- [Section 7, "Documentation Changes"](#)
- [Section 8, "Documentation Accessibility"](#)

1 New Features

This section outlines the new features that are being introduced in the release.

1.1 New Features for 12.1.2.1.2 — Big Data

This section describes the new features for release 12.1.2.1.2.

- **Big data handlers**

The Big Data handlers now supports the Three-Part Table Name and Unified Trail Format features.

- **Token Retrieval Options**

Use the following combinations of new token properties to retrieve these different tokens:

This combination:	Will retrieve these tokens:
<code>gg.includeeggtokens=false</code>	User Tokens
<code>gg.handler.name.includeTokens=true</code>	
<code>gg.includeeggtokens=true</code>	User and GG Tokens
<code>gg.handler.name.includeTokens=true</code>	

This combination:	Will retrieve these tokens:
<code>gg.includeeggtokens=true</code>	No Tokens
<code>gg.handler.name.includeTokens=false</code>	

- **Retrieval of GGS Tokens**

Use the new `gg.includeeggtokens=true` property to retrieve GGS tokens from the data source. When tokens are available, use `op.getToken(<ggstokenname>)` to retrieve the value.

- **Retrieval of All Available Tokens**

When `gg.handler{handlername}.includeTokens=true`, the handler will include *all available user tokens* (key value pairs) to the existing record. This will include the GGS tokens available for the corresponding record.

1.2 New Features for 12.1.2.1.1 — Application Adapters

This section describes the new features for release 12.1.2.1.1.

- **Java User Exit Reporting**

Running the Java user exit in `PASSTHRU` mode disables statistical reporting that allows report counts to be included in the Extract processing report. As of release 12.1.2.1.1, you can trigger the user exit to collect report count statistics and send them to Extract by setting a reporting interval for the newly documented `gg.report.time` property. Then, if Extract is configured to print a report or a request is sent with `SEND REPORT`, the user exit statistics will be included.

- **Simplified Logging**

Logging now defaults to the built-in Java logging, `java.util.logging (JUL)`. Optionally, `java.util.logging` can be changed to either `log4j` or `logback`.

Defaults have been put in place for each of the three possible implementations to simplify the set-up needed for logging. Once the type of logging is identified, the user is not required to enter any properties.

By default, log files are written to the `dirrpt` subdirectory of the Oracle GoldenGate Adapters installation directory. The format of the log file name varies, but all three implementations use the Extract group as part of the default name.

- **Grouped Transactions**

Release 12.1.2.1.1 includes the option to limit transaction size based on the number of operations. See [Section 2.2.1.1, "Transaction Grouping Properties"](#) for details on the properties.

When using the group transactions properties, `gg.handler.name.mode` property and `gg.handler.name.format`, must be set to one transaction per message (transaction mode).

- **Unified Trail Format Support**

Release 12.1.2.1.1 supports the reading of trails that are in the unified trail format. This allows 12.1.x Oracle GoldenGate trails to be processed by the Oracle GoldenGate Adapters flat file writer or user exit. Support for writing unified trail format has been deferred to the next release.

- **Three-Part Table Name Support**

Release 12.1.2.1.1 supports the reading of trails that are written with three-part table names (`database.schema.table`). Both the Java File Writer and Flat File Writer now support three-part table names.

2 New or Changed Properties

This section lists the properties that have been added, changed, or documented for the first time for the release.

2.1 Changed Behavior for 12.1.2.1.4 — Big Data

By default, earlier releases of GoldenGate Big Data Adapter ignored a transaction if an operation failed processing in transaction mode. The entire transaction would be skipped, the adapter process *would not* abend, and the adapter would continue processing the following transactions.

Starting with release 12.1.2.1.4, the adapter process will `ABEND` in the event of an exception in transaction mode to prevent data loss. Fault transactions are fail-fast and will no longer be skipped. This change affects all handlers in transaction mode (operation mode handlers already exhibit this behavior).

Oracle recommends setting the `goldengate.userexit.nochkpt` property to **false** in order to skip operations or transactions that have already been processed and sent to the target. This will avoid having duplicate records.

2.2 New or Changed Properties for 12.1.2.1.1 — Application Adapters

The following properties have been added, changed, or documented for the 12.1.2.1.1 release.

2.2.1 Java Delivery Properties

These new or previously undocumented properties have been implemented for release 12.1.2.1.1.

2.2.1.1 Transaction Grouping Properties These properties specify the minimum and maximum limits for grouping transactions.

- `gg.handler.name.minGroupSize` specifies the minimum number of operations that must be accumulated before a transaction is sent.
- `gg.handler.name.maxGroupSize` specifies the maximum number of operations that are allowed to accumulate before a transaction is sent.

2.2.1.2 Obtain Queue or Topic Object Use `physicalDestination` to obtain the name of the queue or topic object through the `ConnectionFactory` API instead of the JNDI provider.

- `gg.handler.name.physicalDestination=queue_name`

2.2.1.3 Other Java Delivery Properties This section lists other Java Delivery properties that were added or changed for release 12.1.2.1.1.

- `gg.report.time`

`gg.report.time` triggers the user exit to calculate statistics at the indicated interval and sent them to Extract for the processing report. If Extract is configured to print

a report or a request is sent using `SEND REPORT`, these statistics from the user exit are included.

2.2.2 Flat File Writer Properties

The following Flat File Writer property has been changed.

- `goldengate.userexit.utf8mode`

As of release 12.1.2.1.1, the default for `goldengate.userexit.utf8mode` has changed from `FALSE` to `TRUE`.

3 Deprecated Items

This section lists the properties and property options that are deprecated.

3.1 Deprecated Properties for 12.1.2.1.1 — Application Adapters

This section lists the properties and property options that are deprecated for release 12.1.2.1.1.

3.1.1 General Properties

This section lists Adapter properties that are deprecated as of release 12.1.2.1.1.

3.1.1.1 Java Logging Properties The section lists logging properties that are deprecated for release 12.1.2.1.1 Java logging.

- `gg.log.logname`

Specifies the prefix that has the current date appended to it in `yyyymmdd` format to form the log name. The file extension is `log`. On August 3, 2014, an entered value of `writer` will produce a file named `writer_20140803.log`.

- `gg.log.level`

Specifies which levels of messages are reported; errors, warnings, informational, or debug.

- `gg.log.tostdout`

Controls whether log information is written to standard out.

- `gg.log.tofile`

Controls whether log information is written to the log file.

- `gg.log.modules`, `log.level.modules`

Specifies the log level (error, warning, informational, or debug) of the individual source modules that make up the user exit. This is typically used for advanced debugging.

- `gg.log.debug`

Specifies debug of settings for the logging framework.

3.1.1.2 Java Delivery Properties This section lists Java Delivery properties that have been deprecated for 12.1.2.1.1.

- `gg.report.format`

Specifies the template to use for customizing the report format.

4 Installation and Upgrades

This section describes changes to the installation and upgrade procedures.

4.1 Installation Instructions

The installation file downloaded from Oracle is not compressed using gzip; instead it is a standard zip file.

4.1.1 Installation Instruction Changes for 12.1.2.1.1 — Application Adapters

This section lists the changes to the installation instructions for the release.

- The installation instructions for release 12.1.2.1.1 were changed to direct the user first to My Oracle Support (MOS), Patches and Upgrades to see if any patches are available for the requested release. If there are no available patches, the user is directed to <https://edelivery.oracle.com>.
- A new section in the documentation discusses the considerations for running the 12.1.2.1.1 user exit with a non-generic instance of the Oracle GoldenGate core. This includes limitations such as running the adapter with an Oracle GoldenGate instance that is the same major release.
- A new section in the documentation outlines the steps required to install the Java user exit to a non-generic Oracle GoldenGate instance.

4.2 Upgrades

This section lists changes in the upgrade procedure for Oracle GoldenGate Adapters.

4.2.1 From 11.1.x to 12.1.2.1.1

The release 12.1.2.1.1 upgrade process requires additional steps to upgrade the logging configuration.

- When upgrading from 11.1.x to 12.1.2.1.2, logging files may be moved to different directories or may no longer exist. Follow the steps outlined in the *Upgrading the Oracle GoldenGate Adapters* guide to upgrade the logging configuration for 12.1.2.1.1.

5 Corrected Problems

Bug DB *#number* is the number of the Bug DB ticket. For questions on specific tickets or issues, consult Oracle Customer Support.

- [Section 5.2, "Release 12.1.2.1.10 — March 28, 2016 - Application Adapters"](#)
- [Section 5.3, "Release 12.1.2.1.9 — February 23, 2016 - Application Adapters"](#)
- [Section 5.4, "Release 12.1.2.1.8 — January 19, 2016 - Application Adapters"](#)
- [Section 5.5, "Release 12.1.2.1.7 - November 24, 2015 - Application Adapters"](#)
- [Section 5.6, "Release 12.1.2.1.6 — October 27, 2015 — Application Adapters"](#)
- [Section 5.7, "Release 12.1.2.1.5 — October 7, 2015 — Application Adapters"](#)
- [Section 5.8, "Release 12.1.2.1.4 — September 9, 2015 — Big Data"](#)
- [Section 5.9, "Release 12.1.2.1.3 — July 20, 2015"](#)
- [Section 5.10, "Release 12.1.2.1.2 — June 10, 2015 — Big Data"](#)

- [Section 5.11, "Release 12.1.2.1.1 — May 1, 2015 — Application Adapters"](#)
- [Section 5.12, "Release 12.1.2.1.0 — April 9, 2015 — Application Adapters"](#)

5.1 Release 12.1.2.1.11 — September 2, 2016 - Application Adapters

Bug 24416072 - OGG Java Adapter Hadoop not recognize iSeries multi-member table specification
Complete support for native names handling for DB2 for i was added.

Bug 24369897 - Pound Symbol Shows incorrectly when terminal character set is ISO-8859-1 1998
An output encoding property that allows you to control which encoding to use when outputting character data is introduced. This output encoding is specified with the enumeration value of the desired encoding.

5.2 Release 12.1.2.1.10 — March 28, 2016 - Application Adapters

The following tickets have been implemented for the 12.1.2.1.10 Application Adapters.

Bug DB # 2295116 - MAPQA01: Product making inactive @child ord level resulting in a Java error
Access to all user tokens and all Oracle GoldenGate tokens that are consumed by the Oracle GoldenGate Application Adapters implementation was added. This allows Oracle GoldenGate Application Adapters to access these tokens for all platforms.

5.3 Release 12.1.2.1.9 — February 23, 2016 - Application Adapters

The following ticket has been implemented for the 12.1.2.1.9 Application Adapters.

Bug DB # 22518330 - HBASE adaptor key column delimiter/separator is not configurable
The HBase Adapter was modified to add functionality such that the HBase row key delimiter is now configurable. The default value is still a pipe character (|); you can configure the `rowkeyDelimiter` by using the following syntax:

```
gg.handler.name.rowkeyDelimiter=your_delimiter
```

5.4 Release 12.1.2.1.8 — January 19, 2016 - Application Adapters

The following ticket has been implemented for the 12.1.2.1.8 Application Adapters.

Bug DB # 22529284 - AIX: symbol _Z19REP_get_ggs_tkn_bufvlibggjava_ue.so was referenced
AIX contains a limitation such that the adapter module will not be able to process tokens that might be inserted as part of the source trail record.

5.5 Release 12.1.2.1.7 - November 24, 2015 - Application Adapters

The following tickets have been implemented for the 12.1.2.1.7 Application Adapters.

Bug DB #22120256 - This fix preserves the value of transaction indicator flag as seen in the source trail record. The user exit implementation will no longer try to alter the value of transaction indicator flag based to match the downstream or outgoing transaction profile. There will be no logic in the user exit implementation which relies on the value of the transaction indicator flag.

Bug DB #22146222 - Fixed an issue where large source data is truncated to 4000 bytes at the target.

5.6 Release 12.1.2.1.6 — October 27, 2015 — Application Adapters

The following tickets have been implemented for the 12.1.2.1.6 Application Adapters.

Bug DB #21984176 - Fixed an issue where multiple "Begin-Transaction" operations were received for the same source.

5.7 Release 12.1.2.1.5 — October 7, 2015 — Application Adapters

The following tickets have been implemented for the 12.1.2.1.5 Application Adapters.

Bug DB #21836162 - Fixed an issue where the decimal representation of the record position deviates from what is being documented. A record position in decimal representation is (sequence number x 10,000,000,000 + RBA).

5.8 Release 12.1.2.1.4 — September 9, 2015 — Big Data

The following tickets have been implemented for the 12.1.2.1.4 Big Data.

Bug DB #21548335 - Transaction exceptions will no longer be ignored. Instead the adapter process will immediately abend to avoid loss of data. See [Section 2.1, "Changed Behavior for 12.1.2.1.4 — Big Data"](#) for more information.

5.9 Release 12.1.2.1.3 — July 20, 2015

The following tickets have been implemented for the 12.1.2.1.3 Adapters.

Bug DB #21099203 - The v12.1.2 Java Adapter correctly includes tokens in the output, in transaction mode.

5.10 Release 12.1.2.1.2 — June 10, 2015 — Big Data

The following tickets have been implemented for the 12.1.2.1.2 Big Data.

Bug DB #20432399 - GoldenGate Big Data no longer abends with exception while processing Unified Trail Format trail files.

BugDB #21208201 - GoldenGate Adapter correctly reads 12c source definition files.

5.11 Release 12.1.2.1.1 — May 1, 2015 — Application Adapters

The following tickets have been implemented for the 12.1.2.1.1 Adapters.

BugDB #20041455 - Fixed a crash due to incorrect character set handling when `userexit.utf8mode=true`.

BugDB #19982656 - Implemented support to handle `XMLTYPE` columns.

BugDB #19611993 - Fixed a crash in Solaris x86_64 due to incompatible fast compilation done in the Adapter target.

BugDB #13888220 - Implemented support to handle GUID or unique identifier columns.

BugDB #20523498 - The behavior of handler `metacols` transaction indicator (`op.getTxState`), which was changed in 11.2.1.0.3, caused errors in GETENV. To fix this we are deprecating that behavior and reverting to the original behavior.

BugDB #19898067 - The Java rolling file handler has been fixed and works correctly on size, or time, or both.

BugDB #20731562 - Fixed an `IllegalArgumentException` on the `EventObject` source during a rollback event.

BugDB #20657093 - Java user exit now correctly handles different charset data.

BugDB #20639793 - Oracle GoldenGate now issues an exception when the checkpoint by base process is ahead of the one in Java exit.

BugDB #20639752 - Oracle GoldenGate now issues an exception when using grouping with multiple handlers.

BugDB #20525694 - The Java File Writer and Flat File Writer now support the reading of trails that are written with three-part table names (`database.schema.table`).

BugDB #19607132 - The Java handler now supports `physicalDestination` property in order to obtain the queue or topic object from the `ConnectionFactory` API instead of the JNDI provider.

BugDB #18397579 - The Oracle GoldenGate JMS Adapter now correctly encodes characters for all Cyrillic characters.

BugDB #19368122 - Oracle GoldenGate core fixed CLOB chunk retrieval logic, which affected large character object retrieved by the Adapter.

BugDB #17798874 - Oracle GoldenGate core fixed LOB chunk retrieval logic, which affected large object retrieval by the Adapter.

5.12 Release 12.1.2.1.0 — April 9, 2015 — Application Adapters

Initial release.

6 Known Issues

The following are known issues for the release.

6.1 Issues for Release 12.1.2.1.2 — Big Data

The following are known issues for release 12.1.2.1.2 of the Oracle GoldenGate Application Adapters.

- JMS Capture may abend while parsing JMS messages that include three-part table names.

6.2 Issues for Release 12.1.2.1.1 — Application Adapters

The following are known issues for release 12.1.2.1.1 of the Oracle GoldenGate Application Adapters.

- When grouping transactions, the Java user exit may need to wait for additional operations to meet the minimum requirement to send the transaction. As the user exit waits, the Extract checkpoint moves on. If the program abends at this point, it is necessary to manually position the restart in order not to lose data. In this event,

the report file and log file will print a message similar to "Alter Extract to SeqNo *seq_no* and RBA *rba* before restarting the Java Adapter."

- The underlying Extract (OGGCORE 12.1.2.1.1) now has logic to check out-of-order operations based on the sequence id of the captured records. When the process detects the out-of-order condition, it throws the error "ERROR OGG-04036 Positioning sequence ID is out of order." This can happen when the captured records in the JMS queue are coming from two (or more) different data sources, which means each data source will have its own sequence id generation potentially causing interleaved sequence ids. See BugDB #20390986 for more information.
- Tokens within Extract do not work in a Windows environment. As a workaround, you can use Replicat instead of Extract.

7 Documentation Changes

This section discusses changes to the Oracle GoldenGate Adapters documentation.

7.1 Documentation Changes for 12.1.2.1.1 — Application Adapters

Documentation for the Oracle GoldenGate Adapters has the following changes for release 12.1.2.1.1.

- New library pages - the documentation libraries are now organized into a task-based structure designed to make it easier to locate information about the task you are interested in.
- The Flat File and Java Adapter guides have been merged into one *Administering Oracle GoldenGate Application Adapters* guide. This guide is divided into three parts to cover flat file, Java message capture, and Java message delivery.
- Installation and upgrade have been separated out into separate guides.
- Book and chapter titles have been changed to be more task oriented. Each chapter now begins with a brief summary and a list of the topics included in the chapter. Each section begins with a brief summary.
- To improve the documentation for properties:
 - The variables that are part of property names are now included in the documentation. There is also an explanation of the use of variables in property names in section 1.3 of the *Administering* guide.
 - The property descriptions have been changed to include a syntax statement, or an example, or both.
 - The property syntax has been changed to better conform to Oracle standards by removing the curly brackets from variable names and instead denoting variables with italics.
 - An effort has been made to specify the default value for each property variable.
 - Multiple properties that were previously defined together using abbreviated names have been separated and/or changed to use complete names to allow for easier lookups.

8 Documentation Accessibility

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

Access to Oracle Support

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

Oracle GoldenGate Application Adapters Release Notes 12c (12.1.2.1)
E61861-11

Copyright © 2015, 2016, Oracle and/or its affiliates. All rights reserved.

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

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

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

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

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

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

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