

Oracle® Fusion Middleware

Release Notes for Oracle GoldenGate for Big Data

12c (12.3.0.1)

E79606-03

March 2017

Oracle GoldenGate for Big Data 12c streams transactional data into big data systems in real time, raising the quality and timeliness of business insights. This document contains the release notes for the 12c (12.3.0.1.0) release of Oracle GoldenGate for Big Data.

Topics:

[New and Changed Features](#)

[New or Changed Properties](#)

[Upgrade](#)

[Deprecated Items](#)

[Corrected Problems](#)

[Known Issues](#)

[Getting Help with My Oracle Support](#)

New and Changed Features

This section outlines the new features that are being introduced in the release, features that have been enhanced, and changes to existing features.

[Initial Release 12.3.0.1.0 — December 2016](#)

Initial Release 12.3.0.1.0 — December 2016

- The Cassandra and MongoDB Handlers are introduced for use with highly functional NoSQL targets.
- The JDBC Metadata Provider and JDBC Handler are introduced for use with generic JDBC targets. Support for new cloud and data warehousing targets includes Oracle Database, Oracle MySQL, Amazon Redshift, and IBM PureData (Netezza).
- The JSON Formatter was extended to model row data in addition to operation modeled data.
- Support for Coordinated Apply to allow you to scale the throughput by parallelizing processes for your Big Data targets was added.

Note:

Review Understanding What is Supported for more information about support for these new features as well as existing features.

New or Changed Properties

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

[Initial Release 12.3.0.1.0— December 2016](#)

Initial Release 12.3.0.1.0— December 2016

- The HDFS Handler property, `gg.handler.hdfs.openNextFileAtRoll`, was added to support extract, load, transform (ELT) situations. It is only applicable to the HDFS Handler situations that are *not* writing Avro Object Container File (OCF) format or sequence files. When this property is set to `true`, a new HDFS file is created immediately after the prior HDFS file is closed. It allows transformation programs to monitor the output directory and use the appearance of the new file as a trigger to start processing the previous file.
- The following JSON Formatter row modeled properties were added to support this new feature:
 - `gg.handler.name.format`
 - `gg.handler.name.format.flatten`
 - `gg.handler.name.format.flattenDelimiter`
 - `gg.handler.name.format.beforeObjectName`
 - `gg.handler.name.format.afterObjectName`
 - `gg.handler.name.format.pkUpdateHandling`

Upgrade

There are two upgrade paths that you can choose from to upgrade to Oracle GoldenGate for Big Data 12c (12.3.0.1.0) and both are described in *Installing and Upgrading Oracle GoldenGate for Big Data*.

Topics:

Deprecated Items

This section lists all items that were deprecated in each release.

[Initial Release 12.3.0.1.0— December 2016](#)

Initial Release 12.3.0.1.0— December 2016

- The Kafka Handler support is changed such that it no longer supports Apache Kafka 0.8.2.x, Cloudera Distribution of Apache Kafka 1.x (includes Kafka 0.8.2), or Hortonworks HDP Kafka 2.3 (includes Kafka 0.8.2.0).
- Java Delivery using Extract is *not* supported and was deprecated in this release. Support for Java Delivery is only supported using the Replicat process. Replicat provides better performance, better support for checkpointing, and better control of transaction grouping.
- The `goldengate.userexit.nocheckpoint` and `goldengate.userexit.chkptprefix` properties are deprecated.

Corrected Problems

This section describes the problems corrected in each release. The Bug *number* is the number of the BugDB ticket. For questions on specific tickets or issues, consult Oracle Support.

[Patch Release 12.3.0.1.1 — March 2017](#)

[Initial Release 12.3.0.1.0 — December 2016](#)

Patch Release 12.3.0.1.1 — March 2017

Bug 25375346 - Big Data v12.3 output appending invalid values for VARCHAR column

Fixed an issue with extra characters found in the before image of a VARCHAR character column.

Initial Release 12.3.0.1.0 — December 2016

Initial release.

Known Issues

This section describes the known issues, with any available workarounds, identified in each release. The Bug *number* is the number of the BugDB ticket. For questions on specific tickets or issues, consult Oracle Support.

[Release 12.3.0.1.0 — January 2017](#)

[Initial Release 12.3.0.1.0 — December 2016](#)

Release 12.3.0.1.0 — January 2017

Bug 25388733 - Oracle: Collision handling fails on Solaris and AIX platforms for JDBC Replicat

The expected behavior from Replicat is to convert the INSERT operation to an UPDATE operation. On Solaris and AIX platforms this is not happening, instead Replicat pushes the same INSERT operation again and so fails again due to a primary key violation at the target database.

Initial Release 12.3.0.1.0 — December 2016

Amazon JDBC Constraints

BUG 25064525 - JDBC Rep needs schema.tablename quoted in map for Amazon Redshift

If you intend to write and map data to an Amazon Redshift target, you must specify the target object names in lower case and must be double quoted. For example, `"schema"."table"`.

This workaround is due to an Amazon Redshift JDBC driver bug that is tracked under Case 1942771891, <https://console.aws.amazon.com/support/home?region=us-west-2#/case/?displayId=1942771891&language=en>.

Getting Help with My Oracle Support

Use My Oracle Support to find knowledge solutions, workarounds, and other information that is reported by customers, partners, and Oracle employees. My Oracle Support also enables you to open a Service Request. If a patch is required to resolve a service request, you will receive instructions on how to download it from My Oracle Support.

Note:

If you purchased Oracle GoldenGate and support through a distributor, contact your distributor instead of attempting to create a service request through My Oracle Support.

[Viewing the Oracle GoldenGate Knowledge Base](#)

[Creating an Oracle GoldenGate Support Request Ticket](#)

Viewing the Oracle GoldenGate Knowledge Base

To view the Oracle GoldenGate Knowledge Base, follow these steps:

1. Browse to [the My Oracle Support web site at http://support.oracle.com](http://support.oracle.com).
2. Select your language and then log in with your email and Oracle password.
3. Click the **Knowledge** tab.
4. In the **Select a product or product line** field, enter **GoldenGate**, and then select an Oracle GoldenGate product from the context-sensitive list to narrow your results. (You may have to wait a few seconds for this list to appear).
5. In the **Enter search terms** field, enter a search keyword or multiple keywords to focus the query.

Additional information about how to use this tab is obtained by clicking **Help**.

Creating an Oracle GoldenGate Support Request Ticket

If you cannot find an answer to your question or problem in the Knowledge Base, you can open a support request ticket with Oracle Support by following these steps:

1. Browse to [the My Oracle Support web site at http://support.oracle.com](http://support.oracle.com).

2. Select your language, and then sign in with your credentials. If you have not done so, you will have to register to use this web site.
3. Click the **Service Requests** tab.
4. Click **Create SR**.
5. Use the Create Service Request wizard to complete and submit your SR. Additional information about how to use this tab is obtained by clicking **Help**.

Oracle® Fusion Middleware Release Notes for Oracle GoldenGate for Big Data, 12c (12.3.0.1)
E79606-03

Copyright © 2015, 2017, 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.