Oracle® Fusion Middleware

Release Notes for Oracle GoldenGate for Big Data

12c (12.3.2.1.12)
F11079-12
April 2022

Oracle Fusion Middleware Release Notes for Oracle GoldenGate for Big Data, 12c (12.3.2.1.12)

F11079-12

Contents

Preface

Oracle GoldenGate for Big Data 12*c* 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 12*c* (12.3.2.1) release of Oracle GoldenGate for Big Data.

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, such as "From the File menu, select **Save**." Boldface also is used for terms defined in text or in the glossary.
italic *italic*	Italic type indicates placeholder variables for which you supply particular values, such as in the parameter statement: TABLE *table_name*. Italic type also is used for book titles and emphasis.
monospace MONOSPACE	Monospace type indicates code components such as user exits and scripts; the names of files and database objects; URL paths; and input and output text that appears on the screen. Uppercase monospace type is generally used to represent the names of Oracle GoldenGate parameters, commands, and user-configurable functions, as well as SQL commands and keywords.
UPPERCASE	Uppercase in the regular text font indicates the name of a utility unless the name is intended to be a specific case.
{ }	Braces within syntax enclose a set of options that are separated by pipe symbols, one of which must be selected, for example: {*option1* | *option2* | *option3*}.
[ ]	Brackets within syntax indicate an optional element. For example in this syntax, the SAVE clause is optional: CLEANUP REPLICAT *group_name* [, SAVE *count*]. Multiple options within an optional element are separated by a pipe symbol, for example: [*option1* | *option2*].

Related Information

The Oracle GoldenGate Product Documentation Libraries are found at:

https://docs.oracle.com/en/middleware/goldengate/index.html

Additional Oracle GoldenGate information, including best practices, articles, and solutions, is found at:

Oracle GoldenGate A-Team Chronicles

1
Introduction

This chapter introduces the Release Notes for Oracle GoldenGate for Big Data12*c* (12.3.2.1.5).

Topics:

1.1 Latest Release Information

This document is accurate at the time of publication. Oracle will update the release notes periodically after the software release. You can access the latest information and additions to these release notes on the Oracle Technology Network at:

http://www.oracle.com/technetwork/indexes/documentation/index.html

1.2 Purpose of this Document

This document contains the release information for Oracle Fusion Middleware Release for Oracle GoldenGate for Big Data.

Oracle recommends you review its contents before installing, or working with the product.

1.3 System Requirements and Specifications

Oracle GoldenGate follows the Fusion Middleware system requirements and certifications for production environments. For more information, see http://www.oracle.com/technetwork/middleware/ias/downloads/fusion-certification-100350.html.

1.4 Bugs Fixed and Enhancements

This chapter describes the bugs fixed and enhancements at the time of release.

The Bug number is the number of the BugDB ticket. For questions on specific tickets or issues, consult Oracle Support.

1.5 Product Documentation

For complete documentation on Oracle GoldenGate, go to http://docs.oracle.com/goldengate/c1230/gg-winux/index.html.

1.6 Oracle Support

Oracle customers that have purchased support have access to electronic support through My Oracle Support at https://support.oracle.com

1.7 Licensing Information

Licensing information for Oracle Fusion Middleware is available at:

https://shop.oracle.com

Detailed information regarding license compliance for Oracle Fusion Middleware is available at:

Licensing Information

1.8 Downloading and Applying Required Patches

Downloading Patches

Go to My Oracle Support to download the latest software patches.

https://support.oracle.com

See the `README` file in the patch distribution for up-to-date information on the software fixes provided by the patch.

To download and install the latest software patch:

1. Login to My Oracle Support.
2. Click the **Patches & Updates** tab.
3. Under the Patch Search tab, select **Product or Family (Advanced Search)**, and select the **Include all patches in a product family** check box.
4. Enter **Oracle GoldenGate** as the product, select the platform and release, and click **Search**.

The list of currently available patches for Oracle GoldenGate is returned.

For SQL Server CDC Extract you may need these software patches:

- For SQL Server 2012, 2014, and 2016, Microsoft has identified and fixed an issue where some `UPDATE` operations may be written incorrectly to a CDC staging table as an `INSERT` followed by a `DELETE`, rather than a `DELETE/INSERT` pair. This may cause downstream replication issues, such as a primary key violation, therefore Oracle recommends that you apply the Microsoft fix for this issue: https://support.microsoft.com/en-us/help/3030352
- For SQL Server 2016, prior to enabling supplemental logging, ensure that you have patched the SQL Server instance based on the following bug fix from Microsoft: https://support.microsoft.com/en-us/help/3166120/fix-could-not-find-stored-procedure-sys.sp-cdc-parse-captured-column-list-error-in-sql-server-2016 If the instance is not correctly patched with the Microsoft fix, issuing `ADD TRANDATA` against a table for the database may incorrectly report that supplemental logging succeeded when it may not have; therefore no records are captured for that table.

Applying this Patch

You must follow the existing upgrade procedures to overlay the old binaries with the new binaries. In addition, you must rerun `ADD TRANDATA` for each table that is already enabled for `TRANDATA` using these steps:

1. Stop all Oracle GoldenGate processes
2. Follow normal upgrade procedures for binary replacement though do not start any Oracle GoldenGate processes, see Upgrading to Release 12c (12.3.0.1) for Heterogeneous Databases in *Using Oracle GoldenGate for Heterogeneous Databases*.
3. Manually stop the SQL Server CDC Capture job for the database. If the job is processing a large transaction, it may take some time before it actually stops.
4. Ensure that the Extract is stopped.
5. Using GGSCI, run `ADD TRANDATA` again for every table hat you previously enabled it for.

> **Important:**
>
> Do *not* run the `DELETE TRANDATA` command.

6. Manually restart the SQL Server CDC Capture job.
7. Manually restart the Oracle GoldenGate processes (Extract, Replicat, MGR, and so on.)

1.9 Upgrade

There are two upgrade paths that you can choose from to upgrade to Oracle GoldenGate for Big Data 12*c* (12.3.1.1) and both are described in *Installing and Upgrading Oracle GoldenGate for Big Data*.

2
What's New in this Release

This chapter describes the features, enhancements, and changes made to Oracle GoldenGate. Oracle updates the release notes periodically after the software release. This document is accurate at the time of publication.

2.1 12.3.2.1.1 Release — August 2018

- These major features are introduced for the corresponding targets:

  BigQuery Handler
  Use the Google BigQuery Handler, which streams change data capture data from source trail files into Google BigQuery.

  Oracle Cloud Infrastructure Event Handler
  Use the Oracle Cloud Infrastructure Event Handler to load files generated by the File Writer Handler into an Oracle Cloud Infrastructure Object Store.

  Oracle Cloud Infrastructure Classic Event Handler
  Use the Oracle Cloud Infrastructure Classic Event Handler to load files generated by the File Writer Handler into an Oracle Cloud Infrastructure Classic Object Store.

  Connecting to Microsoft Azure Data Lake
  You can connect to Microsoft Azure Data Lake to process big data jobs with Oracle GoldenGate for Big Data.

  Length Delimited Value Formatter
  The Length Delimited Value Formatter is a row-based formatter. It formats database operations from the source trail file into a length delimited value output.

- These changes and enhancements were implemented:

  Cassandra Capture
  Support or SSL configuration and authentication was added.

- The following properties are introduced:

  HBase Handler
  You can omit null fields from being written with `gg.handler.name.omitNullValues=true`.

  Avro Formatters

  - You can write the Avro decimal logical type (per the Avro specification) using `gg.handler.name.format.enableDecimalLogicalType=true`.
  - To handle the Oracle `NUMBER` type, use `gg.handler.name.format.oracleNumberScale=scale`.

- To write the Avro timestamp logical type (per the Avro specification), use `gg.handler.name.format.enableTimestampLogicalType=formatter_spec`.

S3 Event Handler

For Dell ECS, you can set the URL to connect to cloud storage with `gg.eventhandler.name.url`.

All Handlers

- You can consolidate the format of timestamp with this time zone property, `gg.format.timestampWithTimeZone=formatter spec`.
- . For the template name property, `${primaryKey}`, the syntax is extended to allow configuration of key column separator, `${primaryKey[separator_string]}` where the desired separator string is specified within the square brackets].

> **Note:**
>
> Review Understanding What is Supported for more information about support for these new features, as well as existing features.

2.2 Initial Release 12.3.1.1.0 — May 2018

- Theses major features are introduced for the corresponding targets:

Capture for Cassandra

Use Oracle GoldenGate capture (Extract) to get changes from Apache Cassandra databases.

File Writer Handler

The File Writer Handler allows you to write data to a local file system. Additionally, the File Writer employs event handlers to post process data after it is staged to files on the local file system. The event handlers are:

HDFS Event Handler

Uploads staged files to HDFS.

Optimized Row Columnar Event (ORC) Handler

Converts staged files to ORC format and writes either to a local file system or to HDFS.

Parquet Event Handle

Converts staged files to Parquet format and either write to a local file system or to HDFS.

S3 Event Handler

Loads staged files to Amazon S3.

Kafka REST Proxy Handler

Streams change data capture to Kafka via the Confluent Kafka REST Proxy. The Confluent Kafka REST Proxy provides and HTTPS interface for ingest into Kafka.

Oracle NoSQL Handler

Streams change data capture into Oracle NoSQL.

- All Oracle GoldenGate for Big Data Handlers are stateless and only maintain state in the context of the Replicat process that it was running. The File Writer Handler introduces the ability of maintaining state between invocations off the Replicat process.
- The HBase byte fields support was changed so that binary source data is moved into HBase as binary data. It is no longer converted to Base64.
- The `duration` data type was added in Cassandra 3.10. The Cassandra Handler now supports delivery to `duration` data type column. The `duration` value needs to be encoded as a string in your trail.

> **Note:**
>
> Review Understanding What is Supported for more information about support for these new features, as well as existing features.

2.3 Deprecated Items

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

3
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.

3.1 Initial Release 12.3.2.1.0— April 2018

Bug 26856080 — Oracle: Admin Client AUTORESTART help commands are not working

In the Admin Client, the help commands for the `AUTORESTART` parameter are missing.

Workaround
None

Bug 22826452 — Extract does not capture records if WILDCARDRESOLVE IMMEDIATE is used

Do not use `WILDCARDRESOLVE IMMEDIATE` in a multitenant environment if performing DDL on that system. Extract may fail to capture records after the DDL operations.

Workaround
None

Bug 26104564 — Oracle: Deleting AUTORESTART task is killing ER, which was brought up by an AUTORESTART task

When an `execute` task is deleted, the associated `ER` process is killed. This is because the task owns the running state of the ER process.

Workaround
The correct method of creating an “execute” task is as follows:

1. Create the ER process in stopped state.
2. Create the task.
3. Start the task.

Bug 22826452 — Extract does not capture records if WILDCARDRESOLVE IMMEDIATE is used

Do not use `WILDCARDRESOLVE IMMEDIATE` in a multitenant environment if performing DDL on that system. Extract may fail to capture records after the DDL operations.

Workaround
None

Bug 22739872 — Oracle: START MGR can hang for XAGENABLEd setup if the MGR process has abended

When using `XAGENABLE`, if you start manager from `GGSCI` and it immediately abends the command may hang.

Workaround
None

Bug 22967088 — Oracle: IR abends with 3117-ORA-03117: two-task save area overflow

If Integrated Replicat abends with `OGG-00665OCI Error Flushing database inbound server, '..' (status = 3117-: two-task save area overflow)`, disabling `ENABLEMONITORING`.

Workaround
Restart the Replicat.

Bug 19512704 — Oracle: Wrong results with @BINTOHEX and @STRCAT functions

Using `STRCAT` to concatenate a string that has used `BINTOHEX` returns the wrong information.

Workaround
None

Bug 22854712 — Oracle: NOUSERID should force downstream mode

The `NOUSERID` parameter is only supported with downstream Integrated Extract.

Workaround
None

Bug 25043127 — Honor MAPINVISIBLECOLUMNS and NOMAPINVISIBLECOLUMNS in user exit

The `MAPINVISIBLECOLUMNS` and `NOMAPINVISIBLECOLUMNS` parameters do not function properly with user exits.

Workaround
None

XL C++ runtime environment on AIX

You must install the XL C++ runtime environment (RTE) version 13.1 on the AIX systems that you want to use Oracle GoldenGate because it cannot run with older RTE versions.

Workaround
Ensure that the RTE version 13.1 is installed before you install Oracle GoldenGate.

Bug 26564428 — Replicat doesn't map derived objects when schema name is specified with it

When a source schema name is appended to the derived object (indexes and triggers) during DDL (like `CREATE TRIGGER tkggu1.mytrigger1`), then the derived object is not mapped correctly by the Replicat. Since `mapderived` is default for Replicat, the statement should map per the map statements.

Workaround
Do not append a source schema name to the derived object name in DDL statement so that the statement is executed on the correct schema.

Bug 26020817 - Oracle: Extract will have actual LAG + 1 hour during DST transition from DST to Non-DST

Oracle GoldenGate Extract timestamp does not adjust after DST change. An Extract has an actual lag value + 3600 seconds (1 hour) when there is a time transition from DST to non-DST (in November of every year). This recurs during 1 AM to 2 AM non-DST.

Workaround
None

Bug 26812463 - Oracle: GRANT ... TO ... IDENTIFIED BY DDL operation

Combination DDL operations that create a user at the same that as granting permissions to that user are not supported in DDL replication for Classic Extract. For example, `GRANT DBA TO scott IDENTIFIED BY tiger;` is not a supported DDL operation, and Replicat will abend when it encounters these types of operations.

Workaround
You can workaround this by splitting the DDL operation into two separate commands. The first DDL to create the user, and a second DDL to grant the permissions to the new user.

4
Bugs Fixed and Enhancements

This chapter describes the bugs fixed and enhancements at the time of release.

The Bug number is the number of the BugDB ticket. For questions on specific tickets or issues, consult Oracle Support.

4.1 Release 12.3.2.1.12 — April 2022

Upgrade of log4j 2.17.1 to address Security Vulnerability

Oracle recommends that you upgrade to this patch to include these security fixes.

Apache Commons IO Update to at least 2.7

The following security vulnerability has been fixed: **CVE-2021-29425**.

4.2 Release 12.3.2.1.11 — January 2022

Upgrade of log4j to address Security Vulnerability

The following security vulnerabilities have been fixed: **CVE-2021-45046** and **CVE-2021-44228**.

Oracle recommends that you upgrade to this patch to include these security fixes.

After you have applied the patch, do the following:

- Manually remove the older version of log4j related jar files from `ggjava/resources/lib/optional`.
- In the `.properties` file of the replicat or extract groups, add the following: `gg.log.classpath=ggjava/resources/lib/optional/log4j-api-2.16.0.jar:ggjava/resources/lib/optional/log4j-core-2.16.0:ggjava/resources/lib/optional/log4j-slf4j-impl-2.16.0.jar`

4.3 Release 12.3.2.1.9 — July 2020

BUG 31555782 - Oracle GoldenGate for Big Data 12.3.2.1.9

This addresses internal linking library issue for AIX.

4.4 Release 12.3.2.1.8 — June 2020

BUG 31441364 - Generic replicat incurs memory leak when handling before/after image with duplicate maps.

This issue was fixed by correcting issues related to storage of records.

BUG 28065055 - Oracle GoldenGate for BigData Replicat SIGSEGV when processing GG_HEARTBEAT records

This issue was fixed by correcting the `GG_HEARTBEAT` table definition.

4.5 Release 12.3.2.1.7 — May 2020

BUG 31256903 - Duplicate Map - Replication to Kafka Before Values are Empty

This issue was fixed by correcting issues related to before image handling when the same source table is mapped to multiple target tables.

4.6 Release 12.3.2.1.6 — January 2020

Bug 25106122 - Possible SQL injection in Hive automated table creation.

The possible SQL injection in automated Hive table creation functionality was fixed.

30318721 - CVE-2018-11058: ORACLE SECURITY SERVICE UPGRADE

The following security vulnerabilities were fixed:

- CVE-2018-11058
- CVE-2016-0701
- CVE-2016-2183
- CVE-2016-6306
- CVE-2016-8610
- CVE-2018-11054
- CVE-2018-11055
- CVE-2018-11056
- CVE-2018-11057
- CVE-2018-15769

4.7 Release 12.3.2.1.5 — September 2019

BUG 30098309 - gg.schema.sortColumnsStrategy=columnName results in incorrect output values for templates where the ${primaryKeys} keyword is used

The problem of incorrect column indexing when column sorting was enabled using the configuration parameter `gg.schema.sortColumnsStrategy=columnName` was fixed.

BUG 29962268 - BigQuery Handler fails when target schema does not exactly match expected schema.

The functionality was improved to support mapping in target BigQuery tables where the schema does not exactly match the expected schema generated by the BigQuery Handler.

4.8 Release 12.3.2.1.4 — May 2019

Bug 29755778 - Oracle GoldenGate for Big Data 12.3.2.1.0 - 12.3.2.1.3 Initial Kerberos knit for Hadoop and HBase can intermittently fail.

Initial Kerberos authentication (kinit) can intermittently fail for access to Hadoop or HBase. The method called to perform the kinit on the Hadoop `UserGroupInformation` was changed to resolve this.

4.9 Release 12.3.2.1.3 — May 2019

Enh 29385960 - Delimited Text Formatter - Allow configuration of quoting character

The Delimited Text Formatter was enhanced to allow you to configure the quote character that the formatter will use for strings. The `gg.handler.`*`name`*`.format.quoteCharacter` property was added for this support.

Enh 29369088 - The Elasticsearch Handler that feeds data from Oracle GoldenGate to Elasticsearch to send the routing key to route different set of data to different shards in Elasticsearch.

The Elasticsearch Handler was enhanced to allow you to set a template to dynamically resolve the routing key at runtime to control the shard in Elasticsearch to which the message is sent. The `gg.handler.`*`name`*`.routingKeyMappingTemplate` property was added for this support.

Bug 28488882 - OCI Filewriter properly working Replicat shows SEVERE: Error message

Replicat was changed to correctly report the Oracle Cloud Infrastructure Classic Object bucket condition.

Bug 29391606 - Restore removed properties to the JMS Handler.

While changing the JMS Handler, a couple of configuration properties were inadvertently removed, which were added back to the handler.

Bug 29197228 - MongoDB Handler fails to rollback operations when transaction rollback is received

Fixed the MongoDB Handler to rollback operations when a transaction rollback is received.

Bug 29054875 - Output format to JMS topic changes after upgrading to 12.3

The XML Formatter was changed to correctly processes JMS topic changes from previous releases. Three properties were added to support this change: `gg.handler.`*`name`*`.format.missing`, `gg.handler.`*`name`*`.format.missingAfter`, and `gg.handler.`*`name`*`.format.missingBefore`.

Enh 29617633 - S3 Event Handler - Enable the use of the AWS Security Token Service

The S3 Event Handler was enhanced to allow you to access AWS credentials through the Security Token Service if you use federated Single Sign On to access AWS from your on premise environment and you cannot use AWS Identity and Access Management to generate

access keys and secrets. Two properties were added to support this change: `gg.eventhandler.name.enableSTS` and `gg.eventhandler.name.STSRegion`.

Bug 29331365 - Replicat reports message repeatedly causing Replicat to report size drastic increase

A system out print line statement was removed from the ORC Event Handler so that Replicat does not bloat the size of the generated `.rpt` file.

Bug 29281419 - HBase - Error OGG-00453 DDL Replication is not supported for this database.

Replicat was changed to correctly process DDL `MATERIALIZED VIEW` statements.

Bug 27526864— Oracle: Incorrect checkpoint for Coordinated Replicat causing trail files to be purged

An issue when restarting a Coordinated Replicat pointing to an old trail file sequence number was fixed.

Bug 28649552 - Error - OGG-08203 Could not delete DB checkpoint

An erroneous `DBLOGIN` check in `ggsci` when trying to delete a Replicat process was fixed.

Bug 29507343 - Specific HP NonStop data types incorrectly converted in OGGBD.

A mapping issue when handling decimal with decimal point types from HP NonStop was fixed.

Bug 29449724 - JMS Capture AQ - Lag at checkpoint increase

The JMS reader was changed to return an EOF when the read operation doesn't return a new message.

Bug 29548820 - When issued STOP *, all groups go into ABENDED status.

Extract was changed to set the checkpoint ready status to false when processing a SQL operation record, which corrected the erroneous abended status.

Bug 29481919 - Spring Version was changed to 4.3.20 in response to Spring Security Bug CVE-2018-15756.

This security bug is for applications that depend on `spring-webmvc` and `spring-webflux` libraries. Oracle GoldenGate for Big Data has not and does not package either of these libraries nor does it depend on these libraries. It has not and does not utilize the Spring functionality containing the security vulnerability. The change is a formality.

4.10 Release 12.3.2.1.2 — January 2019

Bug 28867246 — Provide regex search and replace functionality for token keys and values

Replicat was changed so that the content replacement regex functionality is applied to token keys and values.

Bug 28754529 — File Writer Handler- Roll on metadata change not working

The File Writer Handler was enhanced so that it rolls files on a metadata change event.

Bug 28716345 — org.apache.avro.SchemaParseException: Illegal character in: RECORD.FILE

An issue with the default field names generated by metacolumns sometimes being illegal in the selected schema was fixed.

Bug 28803633 — No metacolumn support to output the transactional indicator from the source trail file.

The Flat File Writer was changed to add support for a configurable metacolumn to output the transactional indicator from the source trail file. The metacolumn syntax is `${txind}`.

Bug 28819597 — File Writer Handler Replicat crashes server; fails on restart with Java exceptions

The File Writer Handler default buffer size was changed to 1024 bytes (1KB) from 5MB for a smaller memory footprint.

Bug 28476041 — Kafka REST Proxy Handler does not propagate correct schemas

The Kafka REST Proxy Handler was fixed to support sending Kafka Connect messages of different schemas to the same topic.

Bug 28797281 — Metacolumns functionality does not provide support to propagate static values.

A configurable metacolumn was added to the HDFS and Kafka Handlers to propagate a static value. The syntax is `${static[FieldName].StaticValue}`.

Bug 28056315 — Kinesis Handler does not support HTTPS connectivity to proxy server.

Support for HTTPS connectivity to the proxy server for use with the Kinesis Handler and the S3 Event Handler.

Bug 28488865 — ORC Event Handler does not provide support for bloom filter fields.

The ORC Event Handler functionality was modified to add support for bloom filter fields, see ORC Configuring the ORC Event Handler.

Bug 28511829 — S3 Event Handler configuration to write encrypted to S3 bucket

The S3 Event Handler was modified to support configuration for server side encryption. You can choose ASE256 encryption, or select a encryption key or alias from the AWS Key Management System, see S3 Handler Configuration.

Bug 28543382 — File Handler may incorrectly continue processing even when configured event handlers fail initialization

The File Writer Handler was fixed so that if any Event Handler fails initialization, the File Writer Handler abends on initialization.

Bug 28507665 — Regression cause by bug-27804821 where GETUPDATEBEFORES is required

Since the Oracle GoldenGate for Big Data 12.3.1.1.5 release, Replicat does not send before images unless `GETUPDATEBEFORES` is specified. This mode is now the default.

Bug 28065055 — Replicat SIGSEGV when processing GG_HEARTBEAT records

Replicat was changed to handle the mapping of heartbeat tables by automatically creating the target table definition.

Bug 28486402 — EOFDELAYCSECS parameter does not work

The `EOFDELAYCSECS` configuration is not honored and does not have any effect in the actual EOF delay.

Bug 28901887 — WARNING OGG-10173 (rmulhbn.prm) line 41: Parsing error, parameter [map] has unrecognized keyword or extra value "(*ALL)"

Java delivery was changed to return the table name from the record header instead of from the table metadata.

Bug 28395553 — MultiMember table replicating to (*ALL) is not captured

Extract was changed to include definitions for all members of a physical file into the `defs` file if `(*ALL)` is specified.

Bug 28011195 — Port scans cause Extract and Replicat abend with an OGG-1224 Address already in use error

An issue with Extract and Replicat using port scanning software or any other program to connect to an Oracle GoldenGate process on the `localhost` address without using localhost was fixed.

4.11 Release 12.3.2.1.1 — August 2018

Bug 28215602 - Automatically create streams in Kinesis

Two new properties added to the Kinesis Handler to auto create Kinesis streams, `gg.handler.name.enableStreamCreation` and `gg.handler.name.shardCount`.

Bug 28195145 - Initial load process JVM crashes on Linux x86-64

Replicat was changed to handle NULL CSN and XID values.

Bug 28079597 - PURGEOLDEXTRACTS not working correctly

Changed how `PURGEOLDEXTRACTS` operates so that all rules are added to the trail to process when there is no checkpoint in the installation.

Bug 27804821 - Update before image token not showing up in Kafka

The Kafka Handler was changed to use the mapped and target before record to populate the before values.

4.12 Initial Release 12.3.2.1.0 — May 2018

Bug 27525525 - Replicat issues with trail checkpoint and 9-digit trail sequence

Replicat was changed to correct an abend leading to an inconsistent checkpoint file, which caused a problem positioning upon restart. One of the symptoms of this issue is that the Replicat mistakenly assumes that it should reposition to a 6-digit sequence number rather than 9-digit.

Bug 27299518 - SQL is replicated as binary (base64) instead of string

The Kafka Handler was corrected so that a `uniqueidentifier` column is replicated to the target as a base64 string.

Bug 26957243 - CVE-2017-5645: APACHE LOG4J UPGRADE TO 2.9.1

Due to security issues in log4j component, the log4j version was upgraded to log4j 2.9.1.

Bug 26851391 - File Writer Handler - Target data missing Operation counts

The File Writer Handler was changed to count of operations be provided on data files written to HDFS to provide assurance that all records are processed. This is for an ELT scenario.

Enh 25677977 - Enhancement request to correct the DateString format from operation.getTimestamp

Enh 25289618 - Change the HDFS Handler's timestamp format to `yyyy-mm-dd hh:mm¿

Enh 24704644 - Where does "op_ts" milliseconds in time stamp value comes from?

With the HDFS Handler, the timestamp data output format is exactly how it is represented in the trail record. The handler was changed so that you can consolidate all of the timestamp output format using the `gg.format.timestamp property`. The value of this property must follow the specification described in:

https://docs.oracle.com/javase/8/docs/api/java/time/format/DateTimeFormatter.html.

Bug 27012878 — ReverseProxySettings application has two new parameters

The `ReverseProxySettings` application has two new parameters in Oracle GoldenGate version 12.3.0.1 and later:

- `-P`: Password for Service Manager account
- `-u`: Name of Service Manager account to use

Doc 27033479 — New option, SQLMODE, is added for the DBOPTIONS parameter

The `SQLMODE` option is added for the `DBOPTIONS` parameter. When this option is enabled, the `sql_mode` variable is set to `'ANSI_QUOTES'`.

Doc 27061717 — Oracle: Replicat issues replicating TIMESTAMP (6) with TIMEZONE

The daylight saving time adjustment issue when UTC timestamp value converted from local timestamp hits daylight saving time transition period was fixed. From standard time to

daylight saving time 1 hour window. Replicat applies timestamp value 1 hour off from expected value.

Bug 26742447 — SQL Server: Replicat abends with OGG-10124 parsing error for REPERROR value

The `REPERROR` option allowed 1-7 digits of error code only, whereas the SQL Server error code contains 10 digits. To fix this, the regular expression was modified to accumulate 10 digits of SQL Server error code.

Bug 27964981 — Oracle: Extract writes number of bytes > max length for the column to the trail file

An issue with Extract data in the trail file was incorrect when the `VARCHAR` size is more than 32767 bytes was fixed.

Bug 27078084 — Oracle: Integrated Extract captures all columns for delete operations

Fixed as issue where Replicat fails with errors like `OGG-01163 Bad column length (nnnn) specified for colum xxxx`. This issue occurs because the database `NLS_CHARACTERSET` is `AL32UFT8` and `NLS_NCHAR_CHARACTERSET` is `UTF8`. `NVARCHAR2` has a byte limit of 4000 bytes. For `UTF8`, 4000 chars is equal to 4000 bytes. However, Oracle GoldenGate stores `NVARCHAR2` in the trail as `UTF16`, which causes it to expand it to 8000 bytes.

Bug 28043789 — Oracle: @DATE function with @COMPUTE does not work

An issue where the `DATE()` failed if used as an argument of another column mapping function was fixed.

Bug 26869744 — Oracle: Extract Truncating The Column Data When Used with Coltest

An issue with mapping functions truncating numbers that have more digits than the number of bytes needed to store the target number was fixed

Bug 27701581 - Increase the upper limit of the DBOPTIONS parameter options LOBBUFSIZE and XMLBUFSIZE

The upper limit for `LOBBUFSIZE` and `XMLBUFSIZE` was increased from 10485760 to 104857600.