

Qlik Replicate May 2021 - Release Notes



- *In addition to these release notes, customers are also encouraged to read the release notes for all versions later than their current version.*
- *This version includes the Qlik Cloud Landing target endpoint and replication profile, as part of the Qlik Hybrid Data Delivery Early Access Program (Expected GA release: June 2021). For information about the Qlik Hybrid Data Delivery Early Access Program, please contact your Qlik Account Representative or Qlik Support.*

In these release notes:

- *Migration and upgrade (page 2)*
- *What's new? (page 6)*
- *Newly supported endpoint and platform versions (page 14)*
- *End of life/support and deprecated features (page 13)*
- *Resolved issues (page 15)*
- *Known Issues (page 24)*

1 Migration and upgrade

This section describes the issues that you may encounter when upgrading/migrating to the new version.

1.1 Upgrading from an Unsupported Version

If your current Qlik Replicate version is no longer supported, you need to perform two upgrades. First, upgrade to the latest supported version (excluding this one), and then upgrade to this version. If you are unsure what version you need to upgrade to first, contact Qlik Support.

1.2 Microsoft .NET Framework 4.8 Upgrade/Installation Prerequisite

- **Using the Setup Wizard** - It is preferable for .NET Framework 4.8 to be installed on the Replicate Server machine before running Setup. If .NET Framework 4.8 is not present on the machine, Setup will prompt you to install it. This may require the machine to be rebooted when the installation completes.
- **Silent Installation** - The ISS file required for silently installing Replicate must be created on a machine that already has .NET 4.8 installed on it.

1.3 Log Stream Tasks when Upgrading from Replicate 6.4 on Linux

To prevent errors after upgrading, customers with Log Stream tasks running on *Replicate 6.4 for Linux* should perform the following procedure.

For each of your Log Stream tasks:

1. On Replicate 6.4:
 - a. Stop the Log Stream Staging task and make a note of the time.
 - b. For each of the associated Replication tasks, wait for all of the changes from the Log Stream Staging task to be applied to the target endpoints.
To verify this, for each Replication task, switch to **Monitor** view and select the **Change Processing** tab. Once all of the changes have been applied to the target, the **Incoming Changes** gauge and all of the gauges in the **Incoming Changes Details** pane should be empty (and display "0").
 - c. After verifying that all changes have been applied to the target endpoints, stop the Log Stream Replication tasks.
2. Upgrade to Replicate May 2021 and then:
 - a. Select the Log Stream Staging task and open the **Advanced Run Options** window. Select **Date and Time** and specify one minute before the time you stopped the Log Stream Staging task on Replicate 6.4. Click **OK** to resume the task.



Although unlikely, starting the task from a minute before the time it was stopped may result in duplicate records on the target.

- b. Using the same procedure, start all of the associated Replication tasks from the time you started the Log Stream Staging task.

1.4 Kafka Target Endpoint with Hortonworks Schema Registry

To avoid Control Table Namespace conflicts when running multiple tasks, the Control Table Namespace will now be created *without* the task name and schema name.

In light of the above, before upgrading, customers who have configured the Kafka target endpoint to **Publish data schemas to Hortonworks Schema Registry** and who have set the **Schema compatibility mode** to anything other than **None**, need to disable the existing Replicate Control Table subjects in the Hortonworks schema registry.

If needed, you can change the default Control Table Namespace as follows:

1. In Replicate Console, open the task with the Kafka endpoint.
2. In the Task Settings **Message Format** tab, click the **Custom Message Format** button and then click **OK**. Save the task.
3. Export the task using the **Export Task** toolbar button.
4. Open the exported JSON file and add the `control_table_namespace` parameter as follows (replacing `MyNameSpace` with your own value):

```
"task_settings": {
  "source_settings": {
  },
  "target_settings": {
    "queue_settings": {
      "use_custom_message": true,
      "message_shape": {
        "control_table_namespace": "MyNameSpace"
      },
      "use_custom_key": true,
      "key_shape": {
```

5. Save the JSON file and then import it to Replicate using the **Import Task** toolbar button.

In addition, from Replicate April 2020, the default schema compatibility mode for all Control Table subjects will be `None`, regardless of how it is defined in the endpoint settings. Should you wish to use the Schema compatibility mode defined in the Kafka endpoint settings, set the `setNonCompatibilityForControlTables` internal parameter to `false`.

1.5 Deleting the Replicate Self-Signed Certificate after Upgrade

After upgrading, customers that are using Replicate's self-signed certificate (i.e. instead of their own certificate) should perform the following procedure:

1. Delete all *.pem files from <replicate_data_folder>/ssl/data.
2. Restart the **Qlik Replicate Server** service.

This will cause Replicate to generate a new self-signed certificate, thereby resolving any certificate trust issues when connecting to Replicate Console.

Note that if you do not perform the above procedure, the following error will be encountered when connecting to Replicate Console:

```
SYS,GENERAL_EXCEPTION,The underlying connection was closed: Could not establish trust relationship for the SSL/TLS secure channel.
```

1.6 Upgrading an Installation with Multiple Data Folders

When upgrading a Replicate installation with multiple Data folders, only the default Data folder (<Product_Dir>\Data) will be automatically upgraded. The other Data folders need to be updated manually by running the following command:

```
repuctl.exe -d <data_folder_path> setup install
```

1.7 Upgrading SAP Application or SAP Application (DB)

If you are using SAP Application or SAP Application (DB) as a source in a Replicate task, you need to upgrade the SAP transports as follows:

1. Stop all tasks that have a SAP Application or a SAP Application (DB) source endpoint.
2. Upgrade to Replicate May 2021.
3. Upgrade the transports as described in the Replicate Help.
4. Restart the tasks.

1.8 Security hardening: Automatic disabling of the passthrough filter

The passthrough filter allows task designers to control SQL statements executed on source tables during replication. From this version, as part of security hardening, customers will need to explicitly authorize the use of passthrough filters if they wish to continue using them.

After upgrading to this version, any tables with passthrough filters in replication tasks will be suspended and a warning will be issued. If you fully trust the replication task designer, you will then be able to re-enable passthrough filters by setting "enable_passthrough_filter" to "true" in the <product_dir>\bin\repctl.cfg file.

For best security it is recommended to avoid using passthrough filters in replication tasks. If you are unsure about what to do, please contact Qlik Support.

1.9 Google Cloud BigQuery Target - Data Type Changes

From this release, the BYTES and BLOB Replicate data types will be mapped to BYTES (base64) on Google Cloud BigQuery instead of STRING. After upgrading, only new tables will be created with the updated mapping. Existing tables will not be affected unless they are reloaded on the target.

If you wish to continue using STRING instead of BYTES, either define a data type transformation or manually change the data type for the affected target columns post-replication.

1.10 Microsoft Azure SQL Database Target - Duplicate Key Handling Change

In previous versions, when using the Microsoft Azure SQL Database endpoint, duplicate keys would be ignored without issuing an error. Starting from this version, an error will be returned when duplicate keys are encountered.

If you prefer duplicate keys to be ignored (the previous behavior), please contact Qlik Support.

1.11 Qlik Enterprise Manager Compatibility

Qlik Replicate May 2021 is compatible with Qlik Enterprise Manager May 2021 only.

Qlik Replicate May 2021 SR3 is compatible with Qlik Enterprise Manager May 2021 SR3 only.

1.12 New Version and Build Number Format

In previous versions, the build number format for Replicate installation kits was `N.N.N.<build number>` (e.g. `7.0.0.604`). From this version, the following date-based format will be used: `YYYY.MM.<build number>`.

2 What's new?

This section describes the new and enhanced features introduced in Replicate May 2021.

2.1 What's new in endpoints?

Newly Supported Source Endpoints

The following source endpoints are now supported:

- Oracle on Oracle Cloud (via the existing Oracle source endpoint)
- Microsoft Azure Database for PostgreSQL (via the existing PostgreSQL source endpoint)

Newly Supported Target Endpoints

The following target endpoints are now supported:

- Microsoft Azure SQL Managed Instance (via the existing Microsoft SQL Server target endpoint)
- Databricks on Google Cloud



Databricks on Google Cloud is supported from Replicate May 2021 SR3 only.

SAP HANA Log-based CDC

As an alternative to the existing Trigger-based CDC, the SAP HANA source endpoint's new **Log-based CDC** option now enables changes to be captured directly from encrypted or unencrypted logs. When using Log-based CDC, SAP HANA can also be used as a backend database with the SAP Application (DB) source endpoint.

Depending on their environment and corporate security policies, customers can either provide the encryption root keys manually (suitable for rarely changing encryption root keys) or instruct Replicate to retrieve them automatically during runtime (suitable for frequently changing encryption root keys, but requires the ENCRYPTION_ROOT_KEY_ADMIN permission).

SAP Application (DB) source endpoint enhancements

Starting from this version, the following endpoints can now be used with the SAP Application (DB) source endpoint as backend databases:

- IBM DB2 for z/OS source with support for decoding complex pool and cluster tables
- SAP HANA source - Log-based CDC only

Support for Google Cloud BigQuery clustered tables

A **Create tables as clustered by primary key** option has been added to the **Advanced** tab of the Google Cloud BigQuery target endpoint. When this option is selected, the target tables will be created as clustered (according to the first four Primary Key columns that support clustering). In general, clustered tables usually provide significantly faster query performance as well as reducing billing costs.

Kafka target endpoint enhancements

Users can now set a Subject Name Strategy when publishing to Confluent Schema Registry, and access the Confluent Schema Registry via a proxy server.

Subject name strategy support

Kafka endpoint users can now select a subject name strategy when publishing to Confluent Schema Registry.

The following subject name strategies are available:

- Schema and Table Name Strategy - The default.
- Topic Name Strategy

- Record Name Strategy
- Topic-Record Name Strategy



The first strategy (Schema and Table Name Strategy) is a proprietary Qlik strategy while the other three are standard Confluent subject name strategies.

Proxy support

This version introduces support for accessing the Confluent Schema Registry via a proxy server.

Salesforce source endpoint enhancements

The following options have been added to the **Advanced** tab of the Salesforce source endpoint:

- **Use labels for column name** - Enables users to use Salesforce field labels as column names.
- **Replicate encrypted columns** - Enables users to include encrypted columns in the replication task.
- **Perform Full Load using** - You can now choose whether to use **Bulk API** or **SOAP API** for Full Load replication. While Bulk API (the default) is much faster, it is subject to more restrictive quotas and may also sometimes be unavailable due to heavy load.

Microsoft Azure SQL Database - Active Directory authentication

Support for connecting to Microsoft Azure SQL Database using Active Directory authentication has been added.

Support for "Start from timestamp" when using the ODBC with CDC source endpoint

In previous versions, the "Start from timestamp" run option was not supported with the ODBC with CDC source endpoint. From this version, the "Start from timestamp" run option is supported if there is a single context column defined in the Change Processing tab, and its type is **TIMESTAMP**.

Renaming MemSQL target endpoint to SingleStore

To reflect the change to the company name (MemSQL to SingleStore), the MemSQL target endpoint has been renamed to SingleStore.

Handling of Computed Columns when using Microsoft SQL Server-based sources

In previous versions, replication of computed columns from Microsoft SQL Server-based sources (Microsoft SQL Server, Amazon RDS for SQL Server, and Microsoft Azure SQL Managed Instance) was supported in Full Load tasks only. During change processing, any computed columns would be populated with NULL values on the target. This caused issues when the source table column was defined as non-nullable. Consequently, from this version, during change processing, any tables with computed columns will be suspended. If you need to run an Apply Changes and/or Store Changes task that captures changes from tables with computed columns, you should define a transformation to exclude such columns from the task.

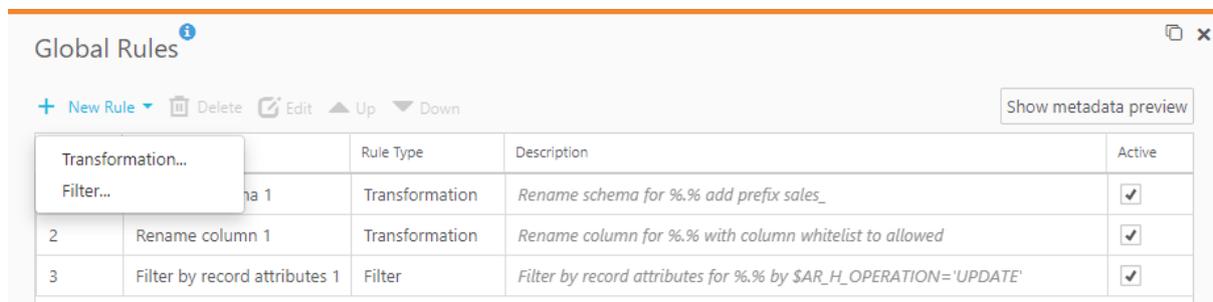
2.2 What's new in engine?

Global rules - transformations and filters

This version introduces significant improvements to the global transformations module. Global transformations allow users to manipulate source data and metadata across multiple tables (in the same task) before it reaches the target. To reflect the new global filtering capabilities, the "Global Transformations" feature has been renamed to "Global Rules". Customers can now define multiple transformation and/or filters that will be executed in their predetermined sequence.

Revamped user interface

The user interface had been redesigned to accommodate the new filtering functionality. In addition to enabling users to define both transformations and filters, users can also set the rule execution sequence (using the Up-Down arrows) and activate/deactivate rules as required.



New global filtering capability

Users can now use the Global Filter Wizard to filter all source records based on column data and/or record attributes. The following filtering options are available:

- **Filter by columns** - Only include records that match specific column data. For example, only include records where Birth Date is later than 02-Feb-2021.
- **Filter by record attributes** - Only include records that match specific record attributes. For example, only include UPDATED or INSERTED records.

New transformation: Replace column value

Use the **Replace column value** transformation to replace the values in the source columns (set in the Transformation scope) with different values in the corresponding target columns.

New metadata variables

The following metadata variables can now be used in global rules:

- \$AR_M_MODIFIED_SCHEMA
- \$AR_M_MODIFIED_TABLE_NAME
- \$AR_M_MODIFIED_COLUMN_NAME
- \$AR_M_MODIFIED_DATATYPE_NAME
- \$AR_M_SOURCE_DATATYPE_LENGTH

- \$AR_M_SOURCE_DATATYPE_PRECISION
- \$AR_M_SOURCE_DATATYPE_SCALE
- \$AR_M_MODIFIED_DATATYPE_LENGTH
- \$AR_M_MODIFIED_DATATYPE_PRECISION
- \$AR_M_MODIFIED_DATATYPE_SCALE

The following data variables can now be used in global rules:

- \$AR_M_SOURCE_COLUMN_DATA
- \$AR_M_MODIFIED_COLUMN_DATA



- "MODIFIED" indicates the object metadata after it has been "transformed".
- The available variables are relevant for the transformation/filter being defined. For example, when defining a **Rename schema** transformation or when filtering by record attributes, only the following variables will be available (new variables indicated in bold): **\$AR_M_MODIFIED_SCHEMA**, **\$AR_M_MODIFIED_TABLE_NAME**, **\$AR_M_SOURCE_TABLE_NAME**, and **\$AR_M_SOURCE_SCHEMA**.

New "replaceChars(X,Y,Z)" function in the Expression Builder

The new `replaceChars(X,Y,Z)` function replaces any character in string `X` that also exists in string `Y` (characters to be replaced) with `Z` (replacement characters) in the same position. This is especially useful for removing non-valid characters from paths and file names.

- If string `Z` (replacement characters) does not include a character that has corresponding position in string `X`, it will be replaced with the first character in string `Z`.
- If string `X` includes a character that does not exist in string `Z`, the original character will be left unchanged.

So, for example, specifying `replaceChars("abcde", "abcd", "123")` would return `1231e`.

Renamed wizard screens

In the Global Transformation Wizard, the following screens have been renamed:

Old Name	New Name
Which Global Transformation?	Transformation Type
What to transform?	Transformation Scope
How to transform?	Transformation Action

New header columns

The following header columns can now be included in transformations:

- **AR_H_DB_COMMIT_TIMESTAMP** - The source database commit timestamp, according to the database server time.



- Relevant for the following source endpoints only: Oracle, Microsoft SQL Server, IBM DB2 for z/OS, Microsoft Azure SQL Managed Instance, and Amazon RDS for SQL Server.
- The `AR_H_DB_COMMIT_TIMESTAMP` header effectively replaces the `use_backend_local_time_in_ct_table_timestamp` internal parameter, which is no longer supported.

- **AR_H_JOB_NAME** - The iSeries job that made the change to the source database.



Relevant for the IBM DB2 for iSeries endpoint only.

- **AR_H_PROGRAM_NAME** - The iSeries program that made the change to the source database.



- Relevant for the IBM DB2 for iSeries endpoint only
- Supported from Replicate May 2021 SR3 only

Enhancements to the "Apply changes using SQL MERGE" option

- Support for the following target endpoints has been added: Snowflake on AWS, Snowflake on Azure, and Snowflake on Google
- In previous version of Replicate, selecting the **Apply changes using SQL MERGE** option in the **Change Processing Tuning** tab would prevent users from being able to define a task-level data error handling policy. From this version, users will be able to define a task-level data error handling policy with the following limitations: The "Ignore record", "Log record to exception table", and "Escalation" options will be unavailable.

Support for LOB Column Replication in UPSERT Error-Handling Mode

In previous Replicate versions, the **Apply Conflicts** "No record found for applying an UPDATE: Insert the missing target record" error-handling option did not support replication of LOB columns (even when the task's **Replicate LOB columns** option was enabled). From this version, Replication of LOB columns when this option is set is now fully supported.

Table reload information in the Change Data Partitions Control Table

In previous versions, when the **Change Data Partitioning** and **Speed partition mode** options were enabled, Replicate would add Full Load partition information to the Change Data Partitions Control Table whenever a table was reloaded (as shown in the image below). From this version, this information will be added to the Change Data Partitions Control Table whenever the **Change Data Partitioning** option is enabled, without needing to enable the **Speed partition mode** option as well.

...	PARTITION_NAME	START_TIME	END_TIME	TABLE_OWNER	TABLE_NAME ...
...	20200804T062000_20200804T062100	2020-08-04 06:20:00.000000000	2020-08-04 06:21:00.000000000	landing	order_details ...
...	20200803T094220_LOAD	2020-08-03 09:42:20.000000000	2020-08-03 09:42:24.617000000	landing	categories ...

Enhanced Kerberos support

In previous versions, customers who wanted to use Kerberos authentication needed to perform tedious manual workarounds to resolve conflicts between Kerberos artifacts installed on their machines and the Kerberos artifacts installed with Replicate. Starting from this version, Replicate is provided with fully functioning Kerberos libraries and utilities, thereby eliminating the need for such workarounds.

Support for build-specific or environment-specific features

This version introduces support for setting build-specific or environment-specific features. As these features are environment-specific, they do not appear as standard options in the user interface. Consequently, they should only be set if explicitly instructed by Qlik Support or product documentation.

These features can be set by clicking **More Options** in the following places:

- The **Advanced** tab of all endpoints
- Task Settings
- Server Settings

Non-nulling of before-image values when using CDC headers

In previous versions, when defining transformations for replication tasks that store changes (in Change Tables or Audit Tables), transformations that leveraged CDC headers (such as using the **User ID** header to prefix a "UID" string to user IDs) would always result in a NULL value in the before-image. The default now is that the before-image values will no longer contain NULL values in such scenarios.

3 End of life/support and deprecated features

This section provides information about End of Life versions, End of Support features, and deprecated features.

3.1 Deprecated endpoint/platform versions

- The following target endpoint versions are no longer supported:
 - Vertica 6.x and 7.x
- The following source endpoint versions are no longer supported:
 - IBM z/OS 2.1 and 2.2

3.2 Deprecated driver versions

To optimize the replication process and minimize connectivity issues, support for old driver versions has been discontinued for several endpoints. For information about which drivers are currently supported for a particular endpoint, refer to the "Prerequisites" section for that endpoint in the Replicate Help.

3.3 Pending deprecation or change

Endpoints

The following endpoint versions will be deprecated in the Replicate November 2021 release:

- Hadoop - Cloudera 5.x
- Hadoop - Hortonworks 2.x
- Hortonworks Data Platform 2.x
- Teradata 15.x
- MySQL 5.6
- IBM DB2 for iSeries 7.2

Change to Kerberos KDC Options - Hadoop and Horton Data Platform Endpoints

From Replicate May 2022, the option to choose **Active Directory** or **MIT KDC** when Replicate is installed on Windows will no longer be available.

Instead, the following options will be available:

- Replicate on Windows - Active Directory KDC only
- Replicate on Linux - MIT KDC only

4 Newly supported endpoint and platform versions

The following source endpoint versions are now supported:

- IBM DB2 for LUW 11.5
- IBM DB2 for iSeries 7.4
- IBM Informix 14.10
- Teradata 17.x
- PostgreSQL 13



PostgreSQL 13 is supported from Replicate May 2021 SR3 only.

The following target endpoint versions are now supported:

- Google Dataproc 2.x
- Teradata 17.x
- PostgreSQL 13



PostgreSQL 13 is supported from Replicate May 2021 SR3 only.

5 Resolved issues

The table below lists the resolved issues for the Replicate May 2021 initial release. For information about issues that were resolved in a subsequent service release, refer to the Replicate May 2021 cumulative release notes ([Qlik_Replicate_May2021_Release_Notes_Accumulated.pdf](#)) provided with the Service Release on the Qlik Download Site.

Resolved Issues

Component/Process	Description	Ref #
IBM DB2 for z/OS Source	When a table was suspended during change processing due to failed parsing, numerous decompression warnings would continue to be reported even though the table was suspended.	N/A
IBM DB2 for z/OS Source	Changes would sometimes not be processed from SAP on IBM DB2 for z/OS tables and the following error would be encountered: DB2z utility (subtype 83) variation 33 (UNIDENTIFIED) at LSN=00D87522D19FE121CA00 was detected for table 'SAPR3'.'MARA'. Operation is ignored (not suspended) based on the endpoint configuration (db2z_endpoint_capture.c:2598) Additional logging was added to assist in troubleshooting the issue.	2053198
IBM DB2 for z/OS Source	When encountering an error caused by the DB2 session being closed by DB2 Manager, the task would fail with a fatal error instead of failing with a recoverable error.	N/A
IBM for DB2 z/OS Source	When encountering an ODBC problem during Change Capture, the task would stop with a fatal error instead of recovering.	N/A
IBM for DB2 z/OS Source	When performing a REORG, the following redundant error would sometimes be encountered (excerpt): DB2z utility (subtype 83) variation 33 (UNIDENTIFIED) at LSN <LSN> was detected for table <name>	N/A
IBM DB2 fro z/OS Source	When the user did not have permission to access the SYSLGRNX table, the task would enter an infinite loop. Now, in such a situation, it will switch to binary search.	N/A
IBM DB2 for z/OS Source	In rare scenarios, the task would fail to start from a timestamp.	2126911

5 Resolved issues

Component/Process	Description	Ref #
IBM DB2 for z/OS Source	When a failure occurred with row decompression, Replicate would retry the task numerous times, resulting in a large backlog of events.	N/A
IBM DB2 for Z/OS Source to Teradata Target	When a source table contained binary columns, applying changes to Teradata in Transactional Apply mode would result in errors when parsing changes on IBM DB2 for Z/OS.	N/A
IBM DB2 for z/OS Source	Added an option to change the DB2z UDTF log severity level while the task is running.	2138502
SAP HANA Source	Records copied from the trigger table to the log table would not be rolled back (from the log table) in the event of a failure to delete the record from the trigger table. Although rare, this would sometimes cause duplicate records on the target.	2069250
SAP HANA source	Capturing changes from tables with primary keys of type VARBINARY would consume excessive memory and take a long time to complete.	2134587
Google Cloud BigQuery Target	The MERGE command would fail for tables with multiple Primary Keys.	N/A
Microsoft SQL Server Source	On rare occasions, when the Alternate backup folder and Replicate had file-level access to the backup log files endpoint settings were configured, Change Processing tasks would enter an infinite loop when capturing changes from a Microsoft SQL Server database that was configured to use FILESTREAM.	N/A
Microsoft SQL Server Source Transformations	When using the AR_H_USER header column to filter transactions based on User ID, the User ID would not be propagated to the column. As this issue only occurs in the customer's environment, additional logging was added to assist in troubleshooting the cause.	2047815
Microsoft SQL Server Source	In rare situations, when parsing a compressed row with a structure that did not correspond to the current table definition, the task would fail.	2094674
Microsoft SQL Server Source	Added support for device type 9 (Azure storage) when the Select virtual backup device types option is enabled in the endpoint settings.	2088559

5 Resolved issues

Component/Process	Description	Ref #
Microsoft SQL Server Source	<p>When the Replicate has file-level access to the backup log files option was enabled, attempting to read a transaction log backup that was restored from Commvault would fail with the following message (shown in verbose logging mode):</p> <pre>[SOURCE_CAPTURE]V: SFMB at offset 0x00000000000010000 (sqlserver_drd_mtf_ map.c:817)</pre>	2084175
Salesforce Source	<p>During Change Processing, changes to tables excluded from the replication task would sometimes prevent the stream position from being updated. This would delay task resumption after recoverable errors.</p>	2070129
Salesforce Source	<p>The connection to Salesforce would reset after a while due to the buffer size being exceeded.</p> <p>The issue was resolved by increasing the buffer size and adding an internal parameter to allow customers to further increase the buffer size if needed.</p>	2074357
SAP Application (DB)	<p>In rare scenarios, when replicating from the BSEG cluster table, the task would fail with the following recoverable error:</p> <pre>[STREAM_COMPONENT]E: Too many 'After Image' entries in BLOB's sequence</pre>	2034536
SAP Application (DB) Source over Oracle 19c	<p>The following error would occasionally be encountered when parsing the online redo logs (excerpt):</p> <pre>The field 'MANDT' doesn't exist in the CDC record for table 'R4S - SALES TRANSACTION'.'LIPS' [1023706] (sapdb_endpoint_ data_record.c:730)</pre> <p>Added logging to try and pinpoint the cause of the error.</p>	2090257
SAP Application (DB)	<p>Full Load of SAP cluster tables without a MANDT column would fail with the following error:</p> <pre>Decompress failed for Cluster table BLOB</pre>	2056366
SAP Application (DB)	<p>Primary Key columns of the SAP pool table would not be parsed correctly when they contained multibyte characters.</p>	2032276
SAP Application (DB) with an Oracle Backend	<p>In rare scenarios, when using Replicate Log Reader to capture changes from wide Advanced Row compressed tables, and the task was configured to perform transformations, the task would sometimes fail.</p>	2157247

5 Resolved issues

Component/Process	Description	Ref #
SAP Application (DB) with an Oracle Backend	When a backend database table of one of SAP tables could not be accessed (for instance, due to insufficient permissions), the task would stop unexpectedly during the Full Load process.	2155526
SAP Application SAP Application (DB) Task Manager	Tasks with a SAP Application (DB) or SAP Application source that were configured as Apply Changes only, would start as Full Load and Apply Changes.	2123964
SAP Application (DB) with Oracle backend	When a recoverable error occurred due to an Oracle connection error, changes would sometimes not appear on the target when the task resumed.	2139031
SAP Application (DB) with an Oracle Backend	When using Replicate Log Reader to capture changes from wide Advanced Row compressed tables, and the task was configured to perform transformations, the task would sometimes fail with the following error (excerpt): <code>The field 'MANDT' doesn't exist in the CDC record</code>	2110294
SAP Application (DB)	When a task is stopped and resumed (or after a recoverable error) and the first transaction that is resent to the target contains cluster table changes, some of the changes in that transaction would sometimes not appear on the target.	2139031
SAP Application (DB)	In rare cases, pool table data columns would be replicated incorrectly.	2128316
Oracle Source - Replicate Log Reader	In rare scenarios, Oracle encrypted columns would be replicated as NULL. Added logging to Oracle Column TDE processing to try and pinpoint the cause of the issue.	2040894
Oracle Source - Replicate Log Reader	In rare situations, DELETE and INSERT operations would not be captured as a result of erroneous parsing of "SPLIT UNDO" operations.	2086367
Oracle Source - Replicate Log Reader	In Oracle 12.1 only, when an OLTP compressed chained row was updated, the affected table would be suspended.	2074822
Oracle Source - Replicate Log Reader	In rare situations, tasks using transformations would sometimes fail on the target due to an issue with the Oracle Deferred Constructor.	2092514
Oracle Source - Replicate Log Reader	In a rarely encountered sequence of redo log events, the task would fail if a transformation was defined for a table with advanced compression.	2103160

5 Resolved issues

Component/Process	Description	Ref #
Oracle Source - Replicate Log Reader	In rare scenarios, several transactions with the same partial transaction ID would be captured as one transaction, resulting in missing UPDATES on the target.	2107559
Oracle Source - Replicate Log Reader	In very rare situations, redo log events would be captured in the wrong order, resulting in inconsistent data on the target.	2135481
Oracle Source	DATE columns would be truncated during Full Load when the bindDateAsBinary internal parameter was set to false.	2118265
Oracle Source	The task would sometimes crash during change capture from a STANDBY environment when the Primary Oracle environment used RAC.	2066583
Security	When using Verbose logging, information about DDLs would sometimes contain user passwords. The problem was resolved by excluding such information from the log.	2096218
Log Stream	When starting a replication task from timestamp, Replicate would sometimes fail to close the log stream file after searching for the start position. This would prevent the log stream task from continuing.	2067415
Log Stream	<p>When using the Log Stream component, the following error would sometimes be encountered:</p> <pre>E: use_backend_local_time_in_ct_table_timestamp is true, but has_source_time_diff is false [1000100] (store_changes.c:1358)</pre> <p>The issue was resolved by adding support for the backend commit timestamp when using Log Stream.</p>	2050262
SAP Extractor Source	When installing the SAP Extractor transport on a new SAP environment, a missing object error would be encountered.	N/A
SAP Extractor Source	When reloading a Full Load only task, the first of the custom extractors would fail to run.	2103568
SAP Extractor Source	<p>Some strings in the target would be missing the last two characters of the original value.</p> <p>In addition, Embedded spaces in string values would be stripped on the target.</p>	2107660
SAP Extractor Source	<p>Certain numeric values on the target would be 1/100th of their value on the source.</p> <p>For example, the value of the 'NETPR' field in SAP would be 4823 while the value in Databricks would be 48.23.</p>	2107657

5 Resolved issues

Component/Process	Description	Ref #
SAP Extractor Source	Due to a failure to clear the delta queue, the same changes would be returned every time the extractor delta process ran.	2106924
SAP Extractor Source	In rare scenarios, locks that were created in the ABAP code would not be released properly.	N/A
SAP Extractor Source	SAP DEC data type columns mapped to REAL8 data type columns would not retain high precision values.	2121616
MySQL Source	When the SSL Mode option was set to Required in the endpoint settings, Replicate would attempt to establish and unsecured connection if a server certificate was not found.	2075672
Amazon RDS for MySQL Source	<p>When replicating from Amazon RDS for MySQL, the task log would contain the following info message numerous times:</p> <pre>00008500: 2021-01-23T01:21:25 [SOURCE_CAPTURE]I: >>> Unsupported or comment DDL: '# Dummy event replacing event type 160 that slave cannot handle. ' (mysql_endpoint_capture.c:1703)</pre> <p>With the fix, the message will only be shown once at info level. At trace level, the message will continue to be shown as the event occurs.</p>	2135640
Microsoft Azure Synapse Analytics Target	<p>Unicode characters would not be replicated correctly into wide columns in Transactional Apply mode.</p> <p>The issue was fixed using an internal parameter.</p>	2098897
Microsoft Azure Synapse Analytics Target Graphical User Interface	It was not possible to set a value in the port field in the Advanced tab.	2141421
IBM DB2 for iSeries	<p>The following redundant warning would sometimes be reported for tables that were not included in the replication task (excerpt):</p> <pre>DROP/RENAME TABLE commands are not currently supported</pre>	2047709

5 Resolved issues

Component/Process	Description	Ref #
Aurora for PostgreSQL Source	When resuming a task, the following fatal error would sometimes occur due to redundant checks: The first begin LSN '00000CCF/F266E830' is higher of stream position LSN '00000CCF/F266D340'. Tables must be reloaded.	2113054
PostgreSQL-based Sources	The WAL slot would constantly grow resulting in degraded performance.	2117811
PostgreSQL-based Sources	Excessive memory consumption would sometimes be encountered during change processing.	2114577
PostgreSQL-based Sources	In rare scenarios, when a large transaction and a small transaction are captured simultaneously, stopping and resuming the task <i>after</i> the small transaction was applied but <i>before</i> the second (large) transaction was processed, would result in missing changes.	2117811
PostgreSQL-based Sources	In rare scenarios, when two transactions occurred while a task was stopped (one small and one large), when the task was resumed, the task would fail with following error: The first event LSN '00000005/9AFA3D38' is higher of stream position LSN '00000000	N/A
PostgreSQL Source	Resuming a task from LSN would sometimes not work.	N/A
PostgreSQL Target	When using the internal parameter <code>psqlReadCommandsFromFile</code> , tables or schemas with capital letters would not be replicated.	1943391
ARC Source Endpoints	Starting a task with an ARC-based endpoint on Linux, would sometimes fail with a recoverable error as the process was unable to determine the correct start position in the CSV file.	2112138
AIS Source	Restarting a task from timestamp would sometimes not be performed correctly with some OS platforms.	2112138
Graphical User Interface	Tasks configured with Character Substitution could not be saved.	2124045
Graphical User Interface	Restored the capability to connect to S3 via a proxy server to the endpoint settings.	2128844

5 Resolved issues

Component/Process	Description	Ref #
Teradata Target	In rare scenarios, the batch optimized apply operation would fail to delete the Replicate net changes table with the following error: NativeError: -2652 Message: [Teradata][ODBC Teradata Driver][Teradata Database](-2652)Operation not allowed	2086486
Microsoft SQL Server Source	The task would sometimes fail when replicating tables with a large number of partitions as a none sysadmin user.	2112203
Microsoft SQL Server Source	The Microsoft SQL Server source endpoint would incorrectly report very high latency.	N/A
Tasks	When stopping or starting a task, there would sometimes be a prolonged delay before the task status changed.	2139808
Engine	Importing a task with a transformation would sometimes set an incorrect table status, causing the task to reload the table unnecessarily.	2128345
Engine (CDC)	The following warning would flood the INFO logging, causing increased latency (excerpt): [INFRASTRUCTURE]W: The transaction timestamp already exists in an earlier partition.	2065627
Engine	When the Metadata only run option was enabled, the parallel load feature would replicate data as well.	2168808
Kafka target	When using the Confluent Schema Registry with ACL authorization, the following error would be encountered when running the task: error code 40301: 'User is denied operation ReadCompatibility on this server The issue was resolved using an internal parameter that bypasses the configs URL call (thereby eliminating the need for super-user permission).	N/A
Kafka Target	Target latency would increase when the last operation captured from the source was a DDL.	N/A
Kafka Target	When processing multiple messages simultaneously, storing the messages in memory while waiting for ACKs from the broker servers would sometimes result in excessive memory consumption on the Replicate Server machine.	N/A

5 Resolved issues

Component/Process	Description	Ref #
Microsoft Azure Synapse Analytics Target	Blob storage folder validation would fail with folders containing upper-case letters.	2155808
Google Cloud BigQuery Target	Loading data into Google Cloud BigQuery would sometimes not complete successfully.	2175799
Oracle Target	In rare scenarios, wide CLOB and CHAR values with non-ASCII characters would not be replicated correctly.	2178291
Microsoft Azure Event Hubs Target	The task would sometimes "detach" from the target with a recoverable error, and take a long time to reconnect.	2155574
AIS Source - IMS	The IMS PSB name could not be used as the User Name in the Replicate header field.	RECOB-2325
Microsoft Azure MySQL Source	The following redundant message would sometimes be returned when replicating columns of type JSON: Col <name> will be skipped since JSON data type is not supported	2161492
IBM DB2 for z/OS Source	When a failure occurred with row decompression, Replicate would retry the task numerous times, resulting in a large backlog of events.	2163452
Databricks Delta Target	When replicating large tables, Full Load would sometimes fail with a query timeout error.	2117837

6 Known Issues

The table below lists the known issues for this release.

Known Issues		
Process	Description	Ref #
Microsoft Azure Synapse Analytics Target	<p>When using Microsoft ODBC Driver 17 on Linux, the following error message is encountered when using the Database name browse button in the endpoint settings (excerpt):</p> <pre>certificate verify failed:unable to get local issuer certificate...Cannot connect to Cloud server</pre> <p>This error occurs because OpenSSL is searching for the Trusted Certificates in the wrong location.</p> <p>Workarounds:</p> <p>The following workarounds are available:</p> <ul style="list-style-type: none"> • Workaround 1: <ul style="list-style-type: none"> • User driver version 13. • Workaround 2: <ol style="list-style-type: none"> 1. Edit the site_arep_login.sh file (default location: <code>/opt/attunity/replicate/bin/</code>) and set the OpenSSL environment variable to point to the correct Trusted Certificates location. <p>Example:</p> <pre>export SSL_CERT_FILE=/etc/pki/tls/cert.pem</pre> <ol style="list-style-type: none"> 2. Restart the Replicate instance by running the following command (shown using the default installation path): <pre>/opt/attunity/replicate/bin/instancename restart</pre> <p>Example:</p> <pre>/opt/attunity/replicate/bin/areplicate restart</pre> 	RECOB-2675
Hadoop High Availability Configuration	<p>When working with a Hadoop High Availability configuration, when a failover occurs during Full Load, the Replicate task does not resume the interrupted load from the Active node.</p>	RECOB-2132

Process	Description	Ref #
Google Cloud BigQuery Target	When a connection error occurs and Replicate recovers the task automatically, the reported number of records replicated during Full Load might sometimes differ from the actual number.	RECOB-2322
Google Cloud BigQuery Target	When the Apply changes using SQL MERGE Change Processing option is enabled, and the value of a source column is changed to NULL in a table without a Primary Key, the NULL value will not be replicated to the target.	RECOB-2780
Microsoft Azure SQL (MS-CDC) Source Endpoint	The Microsoft Azure SQL (MS-CDC) source endpoint is not currently supported, even though it appears in the list of selectable endpoint types.	N/A