

Qlik Replicate November 2020 - Release Notes

This version introduces important security and performance enhancements as well as exciting new features and endpoints.



In addition to these release notes, customers are also encouraged to read the release notes for all versions later than their current version.

In these release notes:

- *Migration and Upgrade (page 2)*
- *New and Enhanced Features (page 9)*
- *Newly Supported Endpoints (page 13)*
- *Newly Supported Endpoint and Platform Versions (page 15)*
- *End of Life/Support and Deprecated Features (page 16)*
- *Resolved Issues and Customer Requested Enhancements (page 18)*
- *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 2020 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 SAP HANA Source - Use Log Table Option

Existing customers who wish to use the new [Use log table](#) option must perform the following procedure beforehand..

1. Stop all tasks replicating from the SAP HANA source endpoint.
2. Back up the “attrep_cdc_changes” table.
3. Edit the following script, and replace the \$SCHEMA\$ with the artifact’s schema:

```
; Create the attrep_log_table

CREATE COLUMN TABLE "$SCHEMA$"."attrep_cdc_log"(

"CHANGE_ROWID" BIGINT CS_FIXED,

"CHANGE_INDX" BIGINT CS_FIXED,

"CHANGE_EVENT_TIME" LONGDATE CS_LONGDATE,

"COPY_TIME" LONGDATE CS_LONGDATE,

"OPER" CHAR(1) CS_FIXEDSTRING,

"TBL_DB_ID" BIGINT CS_FIXED,

"PK_0" NVARCHAR(5000),

"PK_1" NVARCHAR(5000),

"PK_2" NVARCHAR(5000),

"PK_3" NVARCHAR(5000),

"PK_4" NVARCHAR(5000),

"PK_5" NVARCHAR(5000),

"PK_6" NVARCHAR(5000),

"PK_7" NVARCHAR(5000),

"PK_8" NVARCHAR(5000),

"PK_9" NVARCHAR(5000),
```

```
"PK_10" NVARCHAR(5000),

"PK_11" NVARCHAR(5000),

"PK_12" NVARCHAR(5000),

"PK_13" NVARCHAR(5000),

"PK_14" NVARCHAR(5000),

"PK_15" NVARCHAR(5000),

"PK_16" NVARCHAR(5000),

"PK_17" NVARCHAR(5000),

"PK_18" NVARCHAR(5000),

"PK_19" NVARCHAR(5000),

PRIMARY KEY ("CHANGE_ROWID"));

;Copy all changes from attrep_cdc_changes table to the attrep_cdc_log table
INSERT INTO "$SCHEMA$"."attrep_cdc_log"
("CHANGE_ROWID", "CHANGE_INDX", "CHANGE_EVENT_TIME", "COPY_TIME", "OPER", "TBL_DB_ID",
"PK_0", "PK_1", "PK_2", "PK_3", "PK_4", "PK_5", "PK_6", "PK_7", "PK_8", "PK_9", "PK_10", "PK_
11", "PK_12", "PK_13", "PK_14", "PK_15", "PK_16", "PK_17", "PK_18", "PK_19")
(SELECT "$rowid$", "INDX", "EVENT_TIME", CURRENT_TIMESTAMP, "OPER", "TBL_DB_ID",
"PK_0", "PK_1", "PK_2", "PK_3", "PK_4", "PK_5", "PK_6", "PK_7", "PK_8", "PK_9", "PK_10", "PK_
11", "PK_12", "PK_13", "PK_14", "PK_15", "PK_16", "PK_17", "PK_18", "PK_19"
FROM "$SCHEMA$"."attrep_cdc_changes" ORDER BY "$rowid$");

;DELETE all changes from attrep_cdc_changes table
DELETE FROM "$SCHEMA$"."attrep_cdc_changes" WHERE "$rowid$" IN (SELECT "CHANGE_ROWID" FROM
"$SCHEMA$"."attrep_cdc_log");

; COMMIT the changes

COMMIT;
```

4. Run the script one operation at a time with auto-commit off.
The duration of the script job will depend on size of the “attrep_cdc_changes” table.
5. Verify the following:
 - The table “attrep_cdc_log” has been created in the artifact’s schema.
 - The “attrep_cdc_changes” table is empty.
 - The “attrep_cdc_changes” table content has been successfully copied to the new “attrep_cdc_log” table.

1.5 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`.

1.6 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.
```

1.7 Upgrading from Replicate Versions 5.5 with Tasks Configured to use the Microsoft SQL Server Source and/or Target Endpoints

Some of the Microsoft SQL Server endpoint (source and target) internal parameter IDs were changed following the release of Replicate 6.0. To preserve the functionality of these parameters when upgrading from Replicate 5.5, customers should perform the following procedure.

Note that this procedure should only be performed if any of the following internal parameters are set.

Microsoft SQL Server source endpoint internal parameters:

- `accessTlogOnlyModeling`
- `ignoreMsReplicationEnablement`

Microsoft SQL Server target endpoint internal parameter:

- `changeCharParamstoWideSqlType`

Upgrade to the latest version and then:

1. Open the console and delete the relevant internal parameters.
2. Add the deleted internal parameters back again.
3. Save the task and then run it.

1.8 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
```

1.9 Upgrading Replicate Server on Linux

When upgrading from Replicate 5.5 that is configured with the "data" folder in a non-default location, you must add the following parameter to the upgrade command:

```
data=existing_replicate_datadirectory
```

Example (when Replicate is installed in the non-default installation folder):

```
data=/opt/mydatadir/ rpm -U[vh] areplicate-6.5.0-215.x86_64.rpm
```

1.10 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 November 2020.
3. Upgrade the transports as described in the Replicate Help.
4. Restart the tasks.

1.11 Oracle Source Endpoint

Customers upgrading from Replicate 5.5 should set the **Archived redo logs destination identifier** value to the correct DEST_ID. Note that the specified destination must be accessible by Qlik Replicate.

If the **Archived redo logs destination identifier** is not specified (i.e. "0"), Qlik Replicate will use the minimal existing DEST_ID.

1.12 IBM DB2 for z/OS Source Endpoint

Due to enhancements made to the Qlik R4Z (previously known as R4DB2) installation procedure as well as the underlying logic, upgrading the Qlik R4Z component on z/OS from Replicate 5.5 is not supported. Consequently, customers with Replicate 5.5 who wish to use the IBM DB2 for z/OS Source endpoint will need to perform a clean installation of R4Z.

For detailed instructions, refer to the *Qlik R4Z Installation and Configuration Guide*.

1.13 Upgrading the Qlik Replicate Console Only

The following issue applies when upgrading from Replicate 5.5 *only*.

Upgrading only the Qlik Replicate Console in a configuration whereby the Qlik Replicate Console component is installed on one machine and the Qlik Replicate Server component is installed on another, is currently not supported.

Workaround:

1. Uninstall the old Qlik Replicate Console version.
2. Install the new Qlik Replicate Console version.
3. Run the following command from <PRODUCT_DIR>\bin:

```
RepUiCtl.exe repository upgrade --repository ..\Data\GlobalRepo.sqlite
```
4. Restart the Qlik Replicate UI Server service.

1.14 Microsoft SQL Server AlwaysOn Support

Replicate 6.3 introduced improvements to AlwaysOn support that eliminated the need to specify a primary replica. Customers who wish to continue using their existing AlwaysOn configuration after upgrading to Replicate November 2020 from Replicate 5.5 should contact Qlik Support.

1.15 Tasks that were Started from Timestamp



Applies when upgrading from Replicate 5.5.

Usually, when capturing changes from tasks that were started from timestamp, Replicate applies the changes to the target immediately. However, after upgrading, the changes for tasks that were started from timestamp *before* the upgrade will only be applied when the current time is reached, resulting in target latency. The severity of the latency will depend on how far the task has progressed at the time of upgrade. So, for example, the target latency for a task that was started from timestamp one week before the upgrade will be greater than that of a task that was started from timestamp the day before the upgrade (assuming that the amount of changes captured by both tasks is identical).

You can avoid latency issues by waiting until these tasks have reached the current time before upgrading.

1.16 Using External Credentials with an Oracle Endpoint

When using the External Credentials Add-on feature with an Oracle endpoint, extracting the user name from the endpoint settings (as opposed to the Add-on), was the only option. Now, an option to extract both the user name and the password from the compiled Add-on file has been added.

1.17 Qlik Enterprise Manager Compatibility

Qlik Replicate November 2020 is compatible with Qlik Enterprise Manager November 2020 only.

2 New and Enhanced Features

The following section lists the new and enhanced features in this version.

2.1 Security Improvements

Fixed HTTP Header Support (HSTS)

This version of Replicate introduces an option to add fixed HTTP headers to Replicate server responses. This can be used, for example, to add an HSTS header (among others) that might be required by a given organization's security needs. HSTS is a web security policy mechanism that helps to protect websites against man-in-the-middle attacks such as protocol downgrade attacks and cookie hijacking.

Kerberos-SSO Authentication

Kerberos is an enterprise authentication protocol that uses the concept of tickets and three-way authentication to enable users and computers to identify themselves and secure access to resources.

Organizations managing and monitoring their replication tasks via Qlik Enterprise Manager, can now require users to log in to the Qlik Enterprise Manager Console using Kerberos-SSO and administrators can completely externalize and centrally manage users or group memberships using their existing Kerberos infrastructure.

2.2 Performance Improvements

This version introduces significant performance improvements, which are described in detail below.

Speed Partition Mode

By default, Change Data Partitions for all tables in a replication task are registered on the target at the same time. As soon as the partitions are registered, information about them (such as a partition's start and end time) is also published to the Change Data Partitions. This ensures data consistency across all of the replicated tables.

In **Speed partition mode**, rather than waiting for all Change Data files to be uploaded (for all tables), Replicate creates and registers a partition per-table as soon as the partition's first data file is uploaded (to its table).

This is especially useful for consuming applications that need to process changes with a minimum of delay, as the Change Data rapidly becomes available for consumption, even if it is incomplete.

Change Data Partition Cleanup

Over time, Change Data Partitioning can result in a large number of files and partitions accumulating on the target system, which may significantly impact performance. Therefore, to ensure optimal performance, best practice is to delete old partitions from time to time.

To facilitate partition cleanup, Replicate November 2020 introduces a new **Partition Retention** option. The new option is located in the **Change Data Partitioning** section of the **Store Changes Settings** tab (in the **Task Settings** window). When this option is selected (and Change Data Partitioning is enabled), you can set an interval for periodic partition deletion and/or delete partitions on an ad-hoc basis via the new **Delete Old Change Data Partitions** item in the Monitor's **Tools** menu.

Note that to prevent deletion of partitions that are pending consumption, periodic deletion will only be initiated if a retention barrier has been set by the consuming application or applications. The retention barrier must be set using the Enterprise Manager API, which can also be used to perform ad-hoc deletion of old Change Data partitions, and to get a list of all current retention barriers.

For detailed usage instructions, refer to the *Qlik Enterprise Manager API Guide*.



This feature is currently supported with the Microsoft Azure Databricks target endpoint only.

SAP HANA Source Use Log Table Option

The new **Use log table** option significantly reduces latency when capturing changes from SAP HANA. The new option can be used in conjunction with the new **SAP archiving users to filter** option to filter out DELETE operations performed by the specified users. Both of these options can be found in the **Advanced** tab of the SAP HANA endpoint settings.



*To use the **Use log table** option, customers who are currently using the SAP HANA endpoint will need to perform the procedure described in *New and Enhanced Features* (page 9) beforehand.*

IBM DB2 for iSeries - R4I UDTF

By default, changes captured from IBM DB2 for iSeries are filtered on Replicate using the standard display_ journal function. When there is a high volume of changes, this may impact Change Processing performance.

When installed on IBM DB2 for iSeries, the new UDTF significantly improving Change Processing performance by enabling captured changes to be filtered on the source database, instead of on Replicate. To instruct Replicate to use the new UDTF, a new **Enable UDTF capturing** option has been added to the **Advanced** tab of the IBM DB2 for iSeries endpoint settings.

Extended Parallel Load Support

Support for Parallel Load has been extended to the following endpoints:

- Snowflake on Google target
- Google Cloud BigQuery target
- Google Cloud Storage target
- IBM DB2 for z/OS source (by data ranges only)

In Full Load replication mode, you can accelerate the replication of large tables by splitting the table into segments and loading the segments in parallel. Tables can be segmented by data ranges, by partitions, or by sub-partitions.

Extended Change Data Partitioning Support

Support for Change Data Partitioning has been extended to the following target endpoints:

- Google Cloud Storage
- Microsoft Azure ADLS
- Cloudera Data Platform (CDP) Private Cloud
- Databricks on AWS
- Google Data Proc
- Amazon EMR

In a standard replication task, changes are replicated to the target in no particular order. Change Data Partitioning enables processing of Change Data from many tables in a consistent fashion. You can define the duration of partitions as well as the partitioning base time, thereby ensuring overall consistency of the partitioned data.

2.3 Salesforce Source Enhancements

This section describes the Google Cloud BigQuery Target enhancements introduced in this version.

- **Ignore Gap Events** - Salesforce tables with GAP events would be suspended as Replicate does not support such events. An internal parameter was added to ignore GAP events and prevent tables with such events from being suspended.
- **PK Chunking** - Added support for Primary Key chunking of supported tables during Full Load. PK chunking improves performance by splitting bulk queries *on very large tables* into chunks based on the record IDs, or primary keys of the queried records. The new **Enable PK chunking** option is located in the **Advanced** tab of the endpoint settings.

2.4 Google Cloud BigQuery Target Endpoint Enhancements

This section describes the Google Cloud BigQuery Target enhancements introduced in this version.

- In the new **Location** field (located in the endpoint settings' **Advanced** tab), you can now specify where to upload the dataset created by Replicate.
- A new **Default dataset prefix** option has been added to the **Advanced** tab of the endpoint settings. The new option allows you to specify a prefix for the BigQuery dataset where you want the Replicate Control Tables to be created and updated.
- Added support for driver 2.2.4.1011, thereby resolving intermittent connectivity issues.
- **Parallel Load** - Support for [Parallel Load](#) is now available with the Google Cloud BigQuery target endpoint.

- Added support for capturing column DDL operations.
- when source precision > 18 the target will now be created as a number instead of string
- **Applying Changes using SQL MERGE** - In standard operation, the Batch Optimized Apply operation executes separate bulk INSERT, UPDATE and DELETE statements during the replication task. It then executes a single MERGE statement to apply the entire set of NET table changes to the target tables.

While the standard Batch Optimized Apply operation is highly efficient, enabling the **Apply changes using SQL MERGE** option boosts performance even further as it reduces the number of SQL statements run per table from three to one, and significantly lessens I/O. Most UPDATE operations in large, immutable, file-based cloud databases (such as Google Cloud BigQuery), involve rewriting of affected files. With such operations, the reduction of per-table SQL statements from three to one is very significant.



*The **Apply Changes using SQL MERGE** functionality is currently in Beta.*

2.5 Oracle Source Endpoint Enhancements

This version of Replicate introduces the following enhancements to the Oracle endpoint:

- **Oracle Extended Data Types Support** - From Oracle 12.2, support for replication of [extended data types](#) has been added.
- A new **Alternate archived redo logs destination ID** option has been added to the **Advanced** tab of the Oracle source endpoint settings. The new option allows you to specify an alternate archive destination in the event of a failure to read from the primary destination.

2.6 Other Enhancements

- **Table Selection Enhancements:**
 - When selecting tables from source endpoints that do not have a schema/owner, redundant columns will no longer be shown, significantly improving the user experience when selecting tables from such endpoints.
 - A new **Use exact table name** option allows you to exclude and/or include a specific table from the replication task.
- **PostgreSQL Source Data Type** - Support for the JSONB data types was added.
- **Amazon EMR Target** - If you configured your Amazon EMR clusters to use the AWS Glue Data Catalog for metadata storage, you can now tell Replicate to do likewise by selecting the **Store table metadata in AWS Glue Data Catalog** option.
- **Sybase ASE Target** - A nullable BOOLEAN data type will now be mapped to TINYINT in Sybase ASE target.

3 Newly Supported Endpoints

This version of Replicate introduces several new endpoints. For detailed instructions on setting up the endpoints in a Replicate task, please refer to the Replicate Help.

3.1 Microsoft Azure Databricks Delta Target Endpoint

Organizations can now create replication tasks to Microsoft Azure Databricks Delta from any supported source.



*If you wish to use the **Microsoft Azure Databricks Delta** target endpoint, please contact your Qlik Account Manager for a new Replicate license.*

3.2 Cloudera Data Platform (CDP) Private Cloud Target

Organizations can now create replication tasks to Cloudera Private Cloud from any supported source.



- If you wish to use the **Cloudera Data Platform (CDP) Private Cloud** target endpoint, please contact your Qlik Account Manager for a new Replicate license.*
- Customers wishing to migrate replication tasks that currently use the **Hadoop - Cloudera** or **Hadoop - Hortonworks** target endpoints need to create a new task and reload the target and/or contact Product Support for guidance.*

3.3 Oracle Cloud Target

Organizations can now create replication tasks to Oracle Cloud (via the Oracle target endpoint) from any supported source.

3.4 SAP Extractor Source Endpoint

Organizations can now create replication tasks from SAP Extractor to any supported target.

The new endpoint is especially useful for organizations that already have SAP Extractors defined or who need to replicate data directly from the SAP system.



On systems on which the generic Qlik transports are not already installed, installation of the SAP Extractor endpoint transports (R4SAPExtractor transports) may fail with the following error:

Program /QTQVC/EXTR_LAUNCHPAD, Include /QTQVC/EXTR_LAUNCHPAD_TOP: Syntax error in line 000032 Type '/QTQVC/EXTRACTOR' is unknown

If you encounter such an issue, you can either contact Qlik Support for the required transport files, which will allow you to use the SAP Extractor endpoint immediately, or wait for a Service Pack with the fix (due to be released shortly).

3.5 Microsoft Azure Database for MySQL Source Endpoint

Organizations can now create replication tasks from Microsoft Azure Database for MySQL to any supported target.

4 Newly Supported Endpoint and Platform Versions

The following source endpoint versions/editions are now supported:

- Percona 5.7 and 8.0
- MariaDB 10.2-10.5
- PostgreSQL 12

The following target endpoint versions are now supported:

- PostgreSQL 12
- Greenplum 6.x
- Oracle Autonomous Data Warehouse 19.x
- Amazon EMR 6.x
- Kafka is now certified to work with Confluent Platform 6.x

Replicate can now be installed on the following platforms:

- Red Hat 8.x - See the [known issue about Kerberos Authentication](#) below.
- Amazon Linux 2 (on Amazon EC2)

5 End of Life/Support and Deprecated Features

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

5.1 Qlik Replicate 5.5 End of Life

Replicate 5.5 is now End of Life.

5.2 Deprecated Endpoint/Platform Versions

- Windows 2008 R2
- Red Hat 6.x - Qlik products incorporate leading security technologies and modern open standards, so customers can be confident that their data is completely secure. To be able to meet corporate security requirements while also being in a position to rapidly address potential new vulnerabilities, Qlik has decided to focus the development of Replicate Server for Linux on Red Hat 7.x and Red Hat 8.x. As a consequence of this decision, support for Replicate Server for Linux on Red Hat 6.x will be deprecated by the end of the year.
- Red Hat 7.4
- The following target endpoint versions are no longer supported:
 - Kafka target endpoint: 0.8.x, 0.9.x, 0.10.x, and 0.11.x
 - Hadoop Cloudera: Versions 5.8 -5.14
 - Hadoop - Hortonworks: 2.5.x
 - Hadoop - MapR: 5.1
 - Hortonworks Data Platform (HDP): 2.5.x
 - Google Dataproc: 1.2
 - MySQL: 5.5
 - PostgreSQL: 9.3 and 9.4
 - Microsoft SQL Server: 2008 and 2008 R2
 - Oracle: 10.x
 - Teradata: 13 and 14
 - Pivotal Greenplum: 4.2 and 4.3
- The following source endpoint versions are no longer supported:
 - Hadoop Cloudera: Versions 5.5 -5.13
 - Hadoop - Hortonworks: 2.2.x, 2.3.x, and 2.4.x
 - Hadoop - MapR: 4.0, 4.1, 5.0, and 5.1
 - MySQL: 5.5
 - PostgreSQL: 9.4.2 and 9.4.5
 - Microsoft SQL Server: 2008 and 2008 R2
 - Oracle: 10.x
 - IBM DB2 for LUW: 9.7 and 10.1

- IBM DB2 for z/OS: 10
- IBM DB2 for iSeries: 7.1
- Teradata: 13 and 14

5.3 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" topic for that endpoint in the Replicate Help.

5.4 Endpoint/Platform Versions Pending Deprecation

The following platform/endpoint versions will be deprecated in the next version of Replicate:

- Vertica target 6 and 7
- IBM IMS (ARC) source and IBM VSAM Batch (ARC) source on z/OS 2.1 and 2.2

Note also that support for IBM DB2 for z/OS version 10 is planned to be discontinued at the end of 2020.

6 Resolved Issues and Customer Requested Enhancements

The table below lists the resolved issues and customer-requested enhancements for this release.

Resolved Issues and Customer Requested Enhancements

Component/Process	Description	Ref #
Informix Source to File Target	When capturing changes from a table with a serial8 DT column, data for the other target columns may become corrupted when the changes are applied.	2056373
Salesforce source	Full Load of certain tables would sometimes fail due to the Salesforce CSVReader field limit (32,000 bytes) being exceeded.	2047707
SAP HANA source	When the Primary Key GUID was a VARBINARY data type, Change Processing would take a long time.	2049742
SAP Application source with IBM DB2 for LUW backend	The Pool table's VARKEY would be incorrectly parsed, resulting in characters "jumping" to the next column.	2032276
MySQL source to any relational database target.	<p>After stopping a Full Load and Change Processing task defined with a global transformation, removing one of the tables, and then resuming the task, the following error would be encountered for the remaining table(s).</p> <pre>[TARGET_APPLY]E: Failed to build 'where' statement [1022701] (statement_manager.c:1053)</pre> <p>Failed to get update statement for table <name>, stream position.</p>	2038298
PostgreSQL target	<p>Due to an issue with the SSL certificates, all tasks would suddenly stop simultaneously for no apparent reason.</p> <p>The issue was resolved with an internal parameter that adds the SSL parameters to the connection string.</p>	2011220
SAP Application (DB) source	The task would sometimes fail during Change Processing if at least one table did not complete Full Load.	2029030
Microsoft SQL Server source	When the "Alternate backup folder" option was set, Replicate would read the file name incorrectly, resulting in a missing file error.	2007955
IBM DB2 for z/OS source	After upgrading from 6.5, the exclude pattern for selecting tables stopped working.	2030449

6 Resolved Issues and Customer Requested Enhancements

Component/Process	Description	Ref #
PostgreSQL Target	<p>Due to an issue with the SSL certificates, all tasks would suddenly stop simultaneously for no apparent reason.</p> <p>The issue was resolved with an internal parameter that adds the SSL parameters to the connection string.</p>	2011220
Google Cloud BigQuery Target ODBC Target	After encountering a recoverable error, changes for tables without a Primary Key would not be applied on the target.	2015018
Oracle Target	Chinese characters would sometimes not be updated properly, due to an issue with the useUnistrForWideChars internal parameter.	2016391
Replicate General	When Replicate captured a DELETE operation without all the column values (a rare scenario) and then applied the changes to Change Tables, the task would stop abnormally.	1976509
Sap Application (DB) Source	<p>Pool tables would sometimes be parsed incorrectly, resulting in corrupt data on the target and the following error (excerpt):</p> <pre>Failed sending row to table</pre>	1956622
Snowflake Target	Introduced user-settable syntax for improving bulk DELETE operations on Snowflake target.	RPT-25572
IBM DB2 for LUW	<p>When corrupt data was encountered during change processing, the task would enter an infinite loop due to string conversion errors.</p> <p>To resolve this issue, an internal parameter was added which, when set to TRUE, will cause the table to be suspended when corrupt data is encountered.</p>	1975303
SAP HANA Source (Trigger-based)	Capturing changes would sometimes result in excessive memory consumption.	2003744
Microsoft SQL Server Source	<p>When the endpoint's Replicate has file-level access to the backup log files option was enabled, the following error would sometimes be encountered when attempting to capture compressed files:</p> <pre>Encountered an unexpected error</pre>	2001268
IBM DB2 for z/OS Source	In rare scenarios, when the DB2 CLI client returned the RADIX value 10 for the TIMESTAMP column definition, Replicate would fail to parse the columns.	1993890
MongoDB Source	When starting a task from timestamp, the previous changes would not be captured, only the new ones.	1964268

6 Resolved Issues and Customer Requested Enhancements

Component/Process	Description	Ref #
MongoDB Source	When capturing a large amount of cached changes, the changes would sometimes be applied as "normal" events, resulting in unique constraint violations on the target.	1972950
Log Stream	<p>The Log Stream Staging task would sometimes fail with the following error (excerpt):</p> <pre>Failed to read batch from audit file in sub stream</pre> <div style="border: 1px solid gray; padding: 10px; margin-top: 10px;">  <i>After upgrading to this version, customers that encountered such an issue need to start affected tasks from timestamp.</i> </div>	1989737
Microsoft SQL Server Source	<p>When attempting to capture DELETE operations from tables with a non-clustered Primary Key, the following error would be encountered (excerpt):</p> <pre>table - perform_row_lookup(...) failed to look up change data</pre>	1980963
SAP HANA Source	Added a performance trace to the logging in order to identify performance bottlenecks.	1981015
Transformations - Java SDK Endpoints	When adding columns using a table transformation or a global transformation, UPDATE operations would not change the value in the new columns.	1976274
Replicate Server	<p>When trying to save an endpoint with a very long value in one of the fields, the following error would be encountered:</p> <pre>failed to process url...internal server error <1000251></pre>	1975881
Microsoft SQL Server Source	<p>SQL merge operations would sometimes result in missing events on the target.</p> <p>The issue was resolved using an internal parameter, which turns off the demarcation feature by default.</p>	1945604
SAP HANA Source	Added an internal parameter that can be used to split the number of changes being processed into smaller batches. This is useful in situations where a huge number of changes being processed at the same time results in constant timeout errors.	1952418

6 Resolved Issues and Customer Requested Enhancements

Component/Process	Description	Ref #
IBM DB2 for LUW Source	<p>During Change Processing, the following conversion error would be encountered when converting rare Chinese characters (excerpt):</p> <pre>Error converting column 'NAME' in table 'DBTAB01' to UTF8 using codepage 936</pre> <p>The issue was resolved by using a different code page for the conversion.</p>	1943753
Microsoft SQL Server Source	When updates <i>were not</i> performed on VARCHAR(8000) and LOB columns, the data in the corresponding target columns would become corrupted.	1953796
Google BigQuery Target	Resolved disconnection issues by adding support for driver 2.2.4.1011.	1943279
PostgreSQL Source	When the task resumed from a recoverable error, some changes would not be captured.	1908039
Snowflake Target	<p>In "Transactional apply" mode, when a table's Primary Key contained a BINARY data type, UPDATE or DELETE operations would fail with the following error (excerpt):</p> <pre>Cannot convert parameter ''32'' of type [VARCHAR(2)] into expected type [BINARY(32)]</pre>	1951368
SAP HANA Source	<ul style="list-style-type: none"> Added an internal parameter that enables Hints in the fetch events query. Added an internal parameter that can be used to split the number of changes being processed into smaller batches. This is useful in situations where a huge number of changes being processed at the same time results in constant timeout errors. 	1952418
SAP HANA Source	Setting a Log retention period in the endpoint's Advanced tab would have no effect.	1957959
MySQL-to-MySQL Bidirectional	When more than one transaction per second was committed on the source, the timestamp (in the loopback prevention table) would not be changed, resulting in events looping back from the target to the source.	1945480
IBM DB2 for z/OS	The maximum R4Z session duration value has been increased from 20 seconds to 3600 seconds.	1930314
Amazon Kinesis Target	In a Kinesis environment with 200 streams or more, a "Rate Limit Exceeded" error would be encountered when clicking the "Browse" button to select a specific stream.	1931327

6 Resolved Issues and Customer Requested Enhancements

Component/Process	Description	Ref #
Salesforce Source	<p>When a table failed to unload during Full Load, Replicate would erroneously report that Full Load had completed successfully.</p> <p>The issue was resolved by suspending the relevant table(s) and marking the Full Load status as Failed.</p>	1940101
Sap Application with IBM DB2 for LUW backend	<p>When the backend was a non-unicode IBM DB2 for LUW database, the following error would be encountered when starting the task:</p> <pre>RetCode: SQL_ERROR SqlState: 53090 NativeError: -873 Message: [IBM][CLI Driver][DB2/AIX64] SQL0873N Objects encoded with different encoding schemes cannot be referenced in the same SQL statement. SQLSTATE=53090</pre>	193156
Salesforce Source	<p>Unloading of very large tables would sometimes fail. The issue was resolved using an internal parameter to enable automatic primary key chunking for a bulk query job.</p> <p>Supported with the following tables only:</p> <p>Account, Asset, Campaign, CampaignMember, Case, CaseArticle, Contact, Event, EventRelation, Lead, LoginHistory, Opportunity, Task, User, WorkOrder, and WorkOrderLineItem.</p>	1939002
SAP HANA Source	<p>Resuming a task would sometimes result in duplicate and/or missing data on the target.</p>	1925193
Google Cloud Storage Target	<p>Due to network restrictions that prevented connections to accounts.google.com, the endpoint would fail to connect to Google Cloud Storage.</p> <p>The issue was resolved by integrating the endpoint with a newer version of Google API.</p>	1939009
Salesforce Source	<p>The Full Load task could not be stopped manually when an error was encountered.</p>	1940105
Oracle Target Bulk Apply	<p>When applying changes in Batch Optimized Apply mode, missing INSERTs would occur in the following (very rare) scenario:</p> <ul style="list-style-type: none"> • An error would occur during Batch Optimized Apply while capturing changes from two or more tables. • In an attempt to resolve the error, the task was switched to One-by-One apply mode. • In One-by-One mode, errors would also occur, one of which was a table error. 	1920711

6 Resolved Issues and Customer Requested Enhancements

Component/Process	Description	Ref #
Kafka Target Microsoft Azure Event Hubs Amazon Kinesis Data Streams MapR Streams	When the source contained an empty string, messages loaded to the endpoint would sometimes fail with the following error (excerpt): NULL value was unexpectedly encountered	1927270
Microsoft SQL Server Source	When the endpoint was defined to use direct read mode, the task would sometimes stop unexpectedly when reading a table that contained numeric data from a compressed log.	1757012
Change Tables	When the On UPDATE Store after image only option was selected in the Store Changes Settings tab, the Change Table column header__change_mask would always contain the value "00".	1928010
License	A license error would sometimes be encountered with a valid license, resulting in all tasks being stopped.	1914566
SAP Application Source with SAP HANA Trigger-Based Backend	When a DDL for an uncaptured table occurred on the SAP HANA backend, the task would stop abnormally.	1918587
Replicate General	Added support for catching signal types in order to generate a core dump. Supported signals include: SIGSEGV, SIGABRT, SIGALRM, SIGBUS, SIGFPE, SIGHUP, SIGILL, SIGINT, SIGIO, and so on.	1833777
Kafka Target	When using Hortonworks Schema Registry with backward compatibility, specifying a different task name in the Control Table namespace would cause an error.	1906298
Replicate General	The default Endpoint Server would sometimes fail to start, resulting in endpoints developed using the Java SDK not being shown in the UI.	1910726
Oracle Source - Attunity Log Reader	A redo log corruption error was sometimes encountered when running numerous Replicate tasks on Linux.	1897812

7 Known Issues

The table below lists the known issues for this release.

Known Issues

Process	Description	Ref #
Kerberos Authentication	Kerberos authentication is not certified for use with Replicate installed on Linux Red Hat 8.1. Consequently, customers with Replicate on Linux who have endpoints configured with Kerberos authentication should continue using their current Red Hat version until Qlik announces the certification of Kerberos on Linux Red Hat 8.x.	N/A
SAP Extractor	The SAP Extractor endpoint will not be available for selection until the Java connector libraries have been installed and both the Replicate Server service and the Replicate UI Server service have been restarted.	N/A
SAP Extractor Help	<p>In Adding and Managing Source Endpoints > Using SAP Extractor as a Source > Prerequisites, clicking the link outlined below opens the wrong topic.</p> <h3>Prerequisites</h3> <p>The following section describes the prerequisites for working with the Qlik Replicate SAP Extractor endpoint.</p> <ul style="list-style-type: none"> • Install the SAP Java Connector • Install the Qlik Replicate for SAP Client on the SAP machine • Activate the extractors for Replicate <p>The correct page should be:</p> <p>https://help.qlik.com/en-US/replicate/November2020/Content/Replicate/Main/SAPExtractor/extract_Install_Replicate_for_SAP_Client_on_the_SAP_Machine.htm</p>	N/A