

Qlik Replicate Release Notes - April 2020 (v6.6)

This version introduces important features and security enhancements including support for SAML authentication (via Qlik Enterprise Manager), Oracle Autonomous Data Warehouse (an Oracle Cloud Database Service) as a target, Databricks on AWS as a target, publishing data schemas to Confluent Cloud, and replication from MySQL/MariaDB on a Galera Cluster.

Notes 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](#)
- » [New and Enhanced Features](#)
- » [Newly Supported Endpoints](#)
- » [Newly Supported Endpoint and Platform Versions](#)
- » [End of Life/Support and Deprecated Features](#)
- » [Resolved Issues and Customer Requested Enhancements](#)
- » [Known Issues](#)

Migration and Upgrade

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

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.

Kafka Target

Removal of the "Use Kafka Security Settings" Check Box

The **Use Kafka Security Settings** check box was removed from the **Schema Registry Connection Properties** section in the endpoint settings. If the check box was selected before upgrading, then the Kafka endpoint security settings will be propagated to the Confluent/Hortonworks Schema Registry security settings in the the **Schema Registry Connection Properties** section.

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 6.6, 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`.

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.

```

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 and the task and then run it.

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:

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

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
```

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 6.6.

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

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.

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*.

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.

Microsoft SQL Server AlwaysOn Support

Replicate 6.3 introduced improvements to AlwaysOn support that eliminated the need to specify a backup replica. Customers who wish to continue using their existing AlwaysOn configuration after upgrading to 6.6 from Replicate 5.5 should contact Qlik Support.

Tasks that were Started from Timestamp

Note 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.

Qlik Enterprise Manager (formerly known as Attunity Enterprise Manager) Compatibility

Replicate 6.6 is compatible with Qlik Enterprise Manager 6.6 only.

New and Enhanced Features

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

Security Hardening

SAML Authentication

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 SAML. In addition to its single-point-of-authentication and PKI security benefits, SAML also enhances the user experience by allowing users to securely access multiple applications with a single set of credentials entered once.

SAML support has also been added to the Qlik Enterprise Manager public APIs.

Note Qlik Enterprise Manager currently supports SAML 2.0 with the Okta and Microsoft Azure Identity Providers (IdPs).

Kafka Target

The following security enhancements have been made to the Kafka target endpoint:

- » The **Username and Password (SASL/SCRAM-SHA-256)** and **Username and Password (SASL/SCRAM-SHA-512)** authentication types. The **Username and Password** authentication type must be used when [publishing data schemas to Confluent Cloud](#).
- » Using Kerberos authentication when publishing data schemas to Hortonworks Schema Registry

Note Supported with Replicate on Linux only.

Replicate Endpoint Enhancements

SAP HANA Source Enhancements

The following enhancements/changes have been made to the SAP HANA source endpoint:

- » To ensure data consistency when replicating from environments with transactions that routinely exceed 60 seconds, a **Transactional consistency (sec)** field has been

added to the **Advanced** tab.

- » Log-based CDC has been deprecated

Kafka Target Endpoint Enhancements

In addition to the [security enhancements mention above](#), the Kafka target endpoint now supports publishing data schemas to Confluent Cloud.

Oracle Endpoint Enhancements

Oracle Source Endpoint

Replication of Oracle invisible columns is now supported. Replication of invisible columns is enabled by selecting the new **Support invisible columns** option in the **Advanced** tab.

Note Supported from Oracle 12c and with Oracle Instant Client version 12 only.

Oracle Target Endpoint

- » Replication to Oracle Autonomous Data Warehouse (an Oracle Cloud Database Service) is now supported.
- » In an Oracle homogeneous replication task (e.g. Oracle-to-Oracle), if the Primary Key/Unique Index name already exists, Replicate will now generate a new Primary Key/Unique Index name instead of issuing an error.

Amazon RDS for Oracle Source Endpoint

In previous versions, it was only possible to use Oracle LogMiner to access the Oracle redo logs. From this version, it is now possible to use Replicate Log Reader to access the redo logs as well.

Google BigQuery Target Enhancements

The following enhancements have been made to the Google BigQuery target endpoint:

- » Added support for creating and maintaining non-existing datasets on the target.
- » The following limitations *no longer apply*:
 - » Data can only be replicated to a single BigQuery dataset.
 - » The credentials specified in the BigQuery target endpoint will override any existing credentials set by the customer.
 - » All tasks replicating to the BigQuery target endpoint must use the same credentials.

Enhancements to Microsoft Endpoints

Microsoft Azure Synapse Analytics Enhancements (Previously Microsoft Azure SQL Data Warehouse)

The following functionality is now supported:

- » Using Azure Active Directory Authentication as the Authentication method with Replicate on Windows
- » Storing data in ADLS Gen2 storage
- » Connecting via a Proxy Server. The new option can be enabled via the Use Proxy Server check box in the Advanced tab.

Microsoft SQL Server Target Enhancements

You can now use BCP to load tables more efficiently when running Replicate on Linux (by selecting the **Use BCP for loading tables** option in the endpoint's **Advanced** tab).

Microsoft ODBC Driver 17.4.2 Support

Added support for using Microsoft ODBC driver version 17.4.2 with the following endpoints:

- » Microsoft SQL Server source and target
- » Amazon RDS for SQL Server source and target
- » Microsoft Azure Managed Instance source
- » Microsoft Azure SQL Database target
- » Microsoft Azure Synapse Target

IMPORTANT From the next version of Replicate, support for Microsoft SQL Server Native Client 11.0 (Windows) and Microsoft ODBC Driver 13.1.1.0 (Linux) will be discontinued. It is therefore recommended to start working with Microsoft ODBC driver version 17.4.2 as soon as possible.

MySQL-Based Source Endpoint Enhancements

Added support for the following:

- » Replication of JSON data type columns (as the CLOB data type)
- » Replication of XA transactions
- » Galera Cluster for MySQL/MariaDB

IBM DB2 for LUW Source Endpoint Enhancements

Added support for replication from a Purescale environment. This functionality is supported using the **Start processing changes from source change position** Run option only.

IBM DB2 for z/OS Enhancements

Added support for replication of Views.

Note This functionality is supported in Full Load only.

Snowflake on Azure Enhancements

Added an option to use Snowflake as the staging storage type. This option is supported both on Windows and on Linux. Previously, only Blob storage (Windows) was available.

Amazon Redshift Target Enhancements

A **Maximum number of files to copy from Amazon S3 in a single batch** was added to the **Advanced** tab of the endpoint settings. Customers who encounter performance issues may be able to improve performance by adjusting this number.

Newly Supported Endpoints

This section details new and deprecated endpoints/platforms and versions.

Amazon Aurora for PostgreSQL Source

Organizations can now create replication tasks from Amazon Aurora for PostgreSQL to any supported target. For more information on setting up Amazon Aurora for PostgreSQL as a source endpoint in a Replicate task, please refer to the Replicate Help.

Databricks on AWS Target Endpoint

Organizations can now create replication tasks to Databricks on AWS from any supported source. For more information on setting up Databricks on AWS as a target endpoint in a Replicate task, please refer to the Replicate Help.

Newly Supported Endpoint and Platform Versions

The following source endpoint versions are now supported:

- » PostgreSQL 11.4 and 11.5
- » Amazon RDS for PostgreSQL 11.4 and 11.5
- » Microsoft SQL Server 2019
- » Oracle 19.0
- » IBM DB2 for z/OS on IBM z/OS 2.4
- » MariaDB 10.2, 10.3, 10.4, and 10.5

The following target endpoint versions are now supported:

- » PostgreSQL 11.4 and 11.5
- » Amazon RDS for PostgreSQL 11.4 and 11.5
- » Microsoft Azure Database for PostgreSQL 11.5
- » Hadoop Cloudera 6.2 and 6.3
- » Amazon EMR 5.2x
- » Kafka with Confluent Schema Registry version 5.3.1 and Hortonworks Schema Registry 0.8.1
- » Microsoft Azure Databricks 5.x-6.3
- » Microsoft HDInsight 3.6 and 4.0
- » Google Dataproc 1.3 and 1.4

The following platform versions are now supported:

- » Windows Server 2019
- » Red Hat 7.7

End of Life/Support and Deprecated Features

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

Qlik Replicate 5.5 End of Life in November 2020

Replicate 5.5 will be End of Life in November 2020 when Replicate Nov2020 is released.

Endpoint/Platform Versions Pending Deprecation

Platform Versions Pending Deprecation

The following platform versions are pending deprecation:

- » 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 will be deprecated in the next version of Replicate.
- » Windows 2008 R2 will be deprecated in the next version of Replicate.

Endpoint Versions Pending Deprecation

The following endpoint versions are pending deprecation:

- » Support for IBM DB2 for z/OS version 10 is planned to be discontinued at the end of 2020.
- » The following target endpoint versions will be deprecated in the next version of Replicate:
 - » **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 will be deprecated in the next version of Replicate:
- » **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

Resolved Issues and Customer Requested Enhancements

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

Component/Process	Description	Ref #
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
License	A license error would sometimes be encountered with a valid license, resulting in all tasks being stopped.	1914566
Google Dataproc Target	After upgrade, Full Load would sometimes fail with the following error: <code>Failed to find descriptor for '114'</code>	1910952
Google Dataproc Target Amazon EMR Microsoft Azure Databricks Microsoft Azure HDInsight Microsoft Hadoop Hortonworks Data Platform	Data would sometimes be missing from HIVE when the Target storage format was set to Sequence and gzip compression was enabled.	1910511

Component/Process	Description	Ref #
Amazon Kinesis Data Streams Target	<p>When using the endpoint [Browse] button to browse for a specific stream, if more than 100 streams were present on Amazon Kinesis Data Streams, no streams would be listed in the pop-up window and the following error would be returned (excerpt):</p> <pre>Failed to list the available streams on Amazon Kinesis Data Streams</pre> <p>When the stream name was entered manually (instead of using the Browse button), the task would fail.</p>	1907700
License	<p>Introduced additional logging to troubleshoot the issue of tasks sometimes stopping unexpectedly with an invalid license error.</p>	1906840
Kafka Target	<p>Attempting to run more than one replication task using Hortonworks schema registry in "Backward" subject compatibility mode, would fail.</p>	1906298
Kafka Target	<p>UPDATE and DELETE operations performed on source tables with non-nullable LOB columns would fail on Kafka target with the following error (excerpt):</p> <pre>Got unexpected null value in column.</pre>	1905686
Bidirectional	<p>When a bidirectional task was defined, Google Cloud SQL for MySQL target would not be available for selection, even though it is supported.</p>	1903426
Microsoft SQL Server source	<p>In rare situations, when changes were applied to the target slower than they were captured from the source, some of the changes would not be applied to the target.</p>	1901789

Component/Process	Description	Ref #
Microsoft SQL Server Source	During Change Processing, if the endpoint was set with Replicate has file-level access to the backup log files, the task would fail when parsing a rare log operation.	1900340
Notifications	When the Use SSL option was selected, Replicate would fail to send email notifications.	1896572
Amazon Redshift Target	The Primary Key columns would not be recognized as such, which resulted in the endpoint attempting to update a row with the predicate on all the columns. The issue was resolved by using the Primary Key from the source database.	1886464
IBM DB2 for LUW Source	<p>In rare cases, stopping the task while updates are being applied to the target and then resuming the task would result in the following error (excerpt):</p> <pre>Deferred event: BEFORE_UPDATE arrived but prev oper is '0' instead of AFTER_UPDATE</pre>	1885649
Google Cloud BigQuery Target	<p>When working with Google SDK version 274, attempting to create datasets on the target would fail with the following error:</p> <pre>Command failed to load data with exit error code 1, Command output: FATAL Flags parsing error: Unknown command line flag 'dataset'. Did you mean: dataset_id</pre>	1885575
Log Stream with IBM DB2 for z/OS	Replication task would sometimes crash when a DDL was performed on the source table.	1884954

Component/Process	Description	Ref #
Google Dataproc Target	The wrong string size would sometimes be passed to the function when converting from utf16 to utf8, causing the task to crash during Full Load.	1881498
Amazon Kinesis Target	Incorrect record sequence numbers and transaction counters on the target would result in failure of the consuming application.	1878784
SAP Application source	Replicate could not identify Primary Keys in the selected source tables, resulting in inconsistent data on the target.	1878725
Amazon Kinesis Target	Due to the time spent checking for new messages, batches would take several seconds to close, resulting in performance issues. The issue was resolved by streamlining the procedure used to check for new message.	1871866
Replicate Sorter	In very rare cases (e.g. if the task started and then stopped immediately), changes would sometimes not be captured.	1871235
Google Cloud BigQuery Target	When the task was configured to use a schema whose name contained the Google Cloud BigQuery project name, the Drop and Recreate task option would not work during Full Load, resulting in task failure.	1868451
Kafka target	A NULL value being replicated to a non-nullable column would result in task failure instead of a data error.	1867042
IBM DB2 for iSeries Source	If a schema was dropped while the task was stopped, the task would fail instead of continuing with a warning.	1866217

Component/Process	Description	Ref #
Microsoft SQL Server Source	Data from DATETIMEOFFSET columns would be incorrectly converted during Change Processing.	1865060
Server - Scheduler	<p>After upgrading to 6.5, the Scheduled Job "type" would be missing, resulting in the desired action not being performed and the following error:</p> <pre>Global replication task encountered the following error: Unable to lock (write lock) pJobsRwLock.</pre>	1863806
SAP Application (DB) Source	<p>The Replicate Gold Client dictionary name referred to the old SAP tables, resulting in the following error:</p> <pre>Could not find table/view /HTG/GC_TABLES1</pre> <p>The issue was resolved by updating the dictionary table name to refer to the new tables.</p>	1863227
Replicate server	Due to memory not being freed when a connection was closed, Replicate server memory consumption would be unusually high.	1853355
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
IBM DB2 for iSeries Source	In rare cases, the metadata connection would time out during Full Load (after a large table was loaded). This would result in the suspension of all tables that had not been loaded to the target. RPT-22973	1830470

Component/Process	Description	Ref #
SAP Application Source with an IBM DB2 for iSeries or a Microsoft SQL Server backend	The task would stop unexpectedly when capturing DD02L events in the CAPTURE TRACE mode. RPT-23482	1817375
Microsoft SQL Server Source	Capturing a datetime earlier than 01-01-1970 would result in the following error during CDC: <pre> DATETIMEOFFSET value '0001-01-01T00:00:00.0000000' is beyond CDC supported value, which is '1900-01-01T00:00:00.0000000 +00:00 </pre>	1809945
Replicate General	Excessive memory consumption would sometimes be encountered when connecting to Replicate Server.	1804496
Microsoft SQL Server source endpoint	When the endpoint is defined to use direct read mode, the task may sometimes stop unexpectedly when reading a table that contains numeric data from a compressed log.	1757012
Replicate Server	High memory usage was encountered when repeatedly connecting to Replicate Server on Linux.	1755802
Oracle Source - Replicate Log Reader	Capturing changes from an Oracle IOT table without the overflow segment would sometimes result n missing data.	197756
Documentation	A redundant space in the command for configuring the IAM Roles for EC2 Access Type caused an error when the command was copied.	197428

Component/Process	Description	Ref #
PostgreSQL Target	After a lost connection error, the PostgreSQL target would reconnect with the default ODBC PROTOCOL mode which caused the creation of savepoints on the target PostgreSQL.	197419
Oracle Source	When permission for v\$containers was not granted (and not required for the task), the task would fail instead of issuing a warning.	197413
Oracle Source - Replicate Log Reader	The AR_H_USER header on the target would be empty when the following mode was set in the Oracle database: audit_sys_operations=false.	197255
Import	When importing a task from a JSON file to an existing task, the imported task state would appear as "never run" before, regardless of the state of the task it replaced.	197324
Microsoft SQL Server Source	When two tasks with a Microsoft SQL Server source endpoint configured to use the direct-read method were run in parallel, one of the tasks would fail.	197107
SAP HANA Source - Trigger-based CDC	Due to an issue with the SELECT statement, the following error would sometimes occur during change capture (excerpt): <code>General error;339 invalid number: [6930] attribute value is not a number;exception 70006930: attribute value is not a number;</code>	196689
Scheduler	In rare cases, scheduled tasks would start afresh (i.e. with Full Load) instead of resuming.	196136
Oracle Source - General	Tables with a large number of partitions would take a very long time to process.	196054

Component/Process	Description	Ref #
SAP Application Source with SAP HANA backend	When replicating long text, Chinese characters would be displayed as '?' on the target.	195995
Informix Source, Sorter	In some cases, data loss would occur in production due to missing records in the log caused by changes in the stream position format.	195818
Oracle Source - General	When the database included tables with many partitions and sub-partitions, starting the task would take an unusually long time.	195702
Oracle Source - Attunity Log Reader	In rare cases, while running Full Load, the task would fail with the following error: <code>Incorrect numeric value in the column XYZ is changed to 0</code>	195696
Log Stream	When running a task while using Log Stream between the source and target, the table IDs metadata for SAP tables would not specify the same IDs in the Replication task.	195607
UI - Endpoints	When selecting Hortonworks Data Platform (HDP) from the endpoints list, Google Cloud for PostgreSQL would be displayed instead.	195529
Snowflake on AWS	INSERT and UPDATE operations would sometimes not be applied to the target	195393
IBM DB2 for LUW Source	When using Parallel Load, some of the partitions were not loaded to the target.	194940
SAP Application (DB) Source	A maximum of 99 tables would be replicated to the target, even though more than 99 tables were selected for replication	194766

Component/Process	Description	Ref #
SAP HANA Source - Trigger-based CDC	The INDX column in the Replicate cdc_ changes table was defined as LONG instead of BIGINT. This would cause Replicate to stop capturing changes when it reached the limit allowed by the LONG data type. To resolve this issue, the INDX column data type has been changed to BIGINT.	192159

Known Issues

The table below lists the known issues for this release.

Component/Process	Description	Ref #
Google Cloud Storage target endpoint Google Dataproc target endpoint	Accessing these endpoints via an authenticated proxy OR via an HTTPS proxy is currently not supported.	N/A
Salesforce source endpoint	Issues may be encountered when capturing changes from some Salesforce objects with address fields.	RPT-24486
Notifications	When a notification is defined for Full Load completion, notifications will also be sent when "dummy" Full Load completes.	1931817
Amazon RDS for MySQL source endpoint	When running Replicate on Linux, multi-AZ is not supported.	N/A