

Qlik Replicate November 2022 - Release Notes



In addition to these release notes, customers are encouraged to read the IR (Initial Release) release notes for all versions later than their current version as well as the cumulative release notes for the last Service Release of the previous version.

In these release notes:

- *Migration and upgrade (page 2)*
- *What's new? (page 4)*
- *Newly supported versions and third-party software (page 16)*
- *End of life/support features (page 14)*
- *Resolved issues (page 18)*
- *Known Issues (page 54)*

1 Migration and upgrade

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

1.1 Upgrade paths

Qlik supports four Replicate versions (including this one).

Direct upgrade is supported from the last two versions only (including any service releases for those versions). The last two versions are Replicate November 2021 and Replicate May 2022.

Upgrading from Replicate April 2020 (6.6)

If you are upgrading from Replicate April 2020 (6.6), you need to perform the following upgrades:

- a. Upgrade to November 2020 (7.0).
- b. Upgrade to Replicate November 2021.
- c. Upgrade to Replicate November 2022.

Upgrading from unsupported versions

If you are upgrading from Replicate 6.x (starting from 6.3), you need to perform the following upgrades:

- a. Upgrade to Replicate 6.5.
- b. Upgrade to Replicate November 2020 (7.0).
- c. Upgrade to Replicate November 2021.
- d. Upgrade to Replicate November 2022.



If you are upgrading from Replicate 5.5, please 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** - When recording an Install Shield Silent response file (ISS), the recording must be done on a machine with .NET 4.8 already installed on it as the silent installation process does not include a separate installation of .NET framework 4.8.

1.3 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 ensuring that the certificate will be accepted by newer clients (browsers) that select the more secure TLS 1.3 protocol.

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.
```

1.4 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
```

1.5 Compatibility with related products

Qlik Replicate November 2022 is compatible with the following related products only:

- Qlik Enterprise Manager November 2022 or later
- Compose November 2022 or later
- Compose for Data Lakes April 2020 (6.6)

2 What's new?

This section describes the new and enhanced features from Replicate May 2022 Service Release 2 until and including Replicate November 2022 SR1.

2.1 What's new in endpoints?

This section describes the new and enhanced endpoint features.

New features and enhancements in Replicate November 2022 SR1

The following section describes new and enhanced features introduced in Qlik Replicate November 2022 SR1.

Support for 4-byte emoji characters in Snowflake-based endpoints

Support for replication of data with 4-byte emoji characters was added to the following target endpoints:

- Snowflake on AWS
- Snowflake on Azure
- Snowflake on Google

Features and enhancements first introduced in Replicate November 2022 IR (Initial Release)

The following section describes features and enhancements first introduced in Qlik Replicate November 2022 IR.

New Databricks (Cloud Storage) target endpoint

This version consolidates three Databricks target endpoints - AWS, Google Cloud, and Microsoft Azure - into one unified endpoint.

Storage types in the new Databricks (Cloud Storage) endpoint

The screenshot shows a configuration form titled "Cloud Storage Data Access". It has three input fields: "Storage type:", "Bucket name:", and "Bucket region:". The "Storage type:" field has a dropdown menu open, showing five options: "Amazon S3", "Amazon S3", "Google Cloud Storage", "Microsoft Azure Data Lake Storage (ADLS) Gen2", and "Microsoft Azure Blob Storage".

The new Databricks (Cloud Storage) endpoint offers support for:

- SQL Warehouse and All-purpose cluster types
- Four different storage types: Amazon S3, Google Cloud Storage, Microsoft Azure Data Lake Storage (ADLS) Gen2, and Microsoft Azure Blob Storage.

- Databricks implementations on Microsoft Azure, AWS, and Google Cloud Platform.
- Parquet format when using an SQL Warehouse cluster

Customers using the existing Databricks endpoints do not need to take any action as the new endpoint will automatically replace their existing endpoint, after upgrading to Replicate November 2022.

New SAP ODP (Operational Data Provisioning) source endpoint

This version introduces support for SAP ODP (Operational Data Provisioning) as a source endpoint, enabling replication of CDS Views and various SAP BW objects.

New Amazon MSK target endpoint

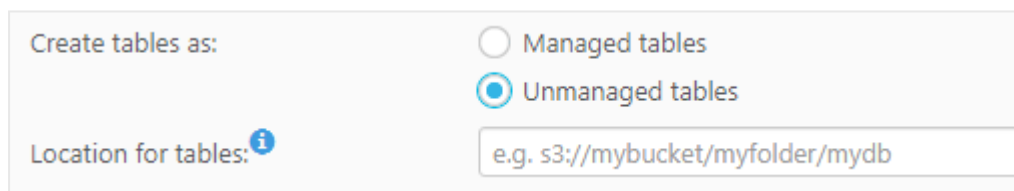
This version introduces support for Amazon MSK as a target endpoint.

Databricks Lakehouse (Delta) enhancements

For customers upgrading from Replicate May 2022 Service Release 2, the following enhancements have been added.

Option to create tables as managed or unmanaged

A new option that allows you to choose whether the target tables will be created as **Managed** or **Unmanaged** has been added to the Databricks Lakehouse (Delta) target endpoint.



Create tables as:

Managed tables

Unmanaged tables

Location for tables: ⓘ

e.g. s3://mybucket/myfolder/mydb

For more information on managed versus unmanaged tables, see <https://docs.databricks.com/lakehouse/data-objects.html>

Apply changes using SQL MERGE

Databricks Lakehouse (Delta) now supports the **Apply changes using SQL MERGE** change processing option. For more information on this option, see https://help.qlik.com/en-US/replicate/Content/Global_Common/Content/SharedEMReplicate/Customize%20Tasks/tasks_applychangtunestab.htm#anchor-4.

MongoDB source endpoint enhancements

MongoDB Atlas

The MongoDB source endpoint now supports replication from MongoDB Atlas, in addition to the standard MongoDB support.

LDAP authentication

Added support for LDAP authentication when using MongoDB Standard.


CDC support when replicating from Google Cloud SQL for PostgreSQL

Support for CDC has been added to the Google Cloud SQL for PostgreSQL source endpoint.

Microsoft Azure Synapse Analytics target endpoint enhancements

Support for the Synapse COPY statement

Customers can now choose whether to load data into Microsoft Azure Synapse Analytics tables using PolyBase or the COPY statement. In certain environments, the COPY statement might offer better performance.



Performance tuning

The new **Use compression** and **Number of files to apply in a batch** (for Full Load) options can improve performance, depending on the size and number of tables being replicated.

Google Cloud BigQuery enhancements

The following options were added to the **Advanced** tab of the Google Cloud BigQuery target endpoint:

- **Proxy support:** Customers can now connect to Google Cloud BigQuery target via a proxy server.
- **Use compression:** When selected, the endpoint will compress (using gzip) the CSV files before uploading them to Google Cloud BigQuery
- **Other:** Customers can now choose **Other** from the **Location** drop-down list and set the location code in the **Region code** field. Region codes are listed in <https://cloud.google.com/bigquery/docs/locations>.

Support for Azure Database for PostgreSQL - Flexible Server as a source

Customers can now replicate from Azure Database for PostgreSQL - Flexible Server using the standard PostgreSQL source endpoint.







Support for IBM Netezza Performance Server (Cloud) as a target

Customers can now replicate to IBM Netezza Performance Server using the standard IBM Netezza target endpoint.

Salesforce (Incremental Load) source endpoint changes and enhancements

Option to select the API automatically

For both Full Load and CDC loading, a **Select API automatically** option has been added. When selected, If the number of records per table matches or exceeds the value specified in the **API switchover threshold (records)** field, Replicate will retrieve the data using the Bulk API (if the object supports Bulk API); otherwise, it will retrieve the data using the SOAP API.

Full Load method: 	Select API automatically 
API switchover threshold (records):	2000  (1 - 100,000,000)
CDC method: 	Select API automatically 
API switchover threshold (records):	2000  (1 - 100,000,000)

Relocation of the "Check for changes" option

The **Check for changes** option has been moved to the **Advanced** tab.

Support for Azure Active Directory Service Principal and Azure Active Directory authentication

Support for Azure Active Directory Service Principal authentication was added to the following endpoints:

- Microsoft Azure SQL Managed Instance source
- Microsoft Azure SQL (MS-CDC) source
- Microsoft Azure SQL Database target

Support for Azure Active Directory authentication was added to the following endpoints:

- Microsoft Azure Managed SQL Instance source
- Microsoft Azure SQL (MS-CDC) source



The Azure Active Directory Service Principal authentication method requires ODBC driver 18.1.

Support for triggered events when using the Microsoft Azure ADLS target endpoint

When the **Storage type** is set to **Azure Data Lake Storage (ADLS) Gen2**, the Microsoft Azure ADLS target endpoint now supports Azure Storage Events. To learn more about Azure Storage Events, see <https://learn.microsoft.com/en-us/azure/storage/blobs/storage-blob-event-overview>

Features and enhancements first introduced in Replicate May 2022 SR2

The following section describes features and enhancements first introduced in Replicate May 2022 SR2.

New Google Cloud SQL for SQL Server target endpoint

Customers can now replicate data from any supported source to **Google Cloud SQL for SQL Server**.

New Microsoft SQL Server (MS-CDC) source endpoint

Customers can now replicate data from Microsoft SQL Server using the MS-CDC mechanism. When replicating data from Microsoft SQL Server using MS-CDC, changes are captured from Microsoft Change Tables as opposed to the SQL Server transaction log. Using MS-CDC is simpler from a setup perspective and it eliminates the need to use a high-privileged account to access the transaction logs. On the other hand, because the MS-

CDC mechanism is based on Change Tables, the replicated transaction patterns are different from the original DML transaction that occurred in the database. It should be noted that while this is still transactionally consistent, it does not represent the original transactions.

New Databricks Lakehouse (Delta) target endpoint

This version introduces the new Databricks Lakehouse (Delta) target endpoint, offering support for:

- Three different storage options: Amazon, Google, and Microsoft Azure Data Lake Storage Gen2.
- Databricks implementations on Azure, AWS, and Google Cloud Platform.
- Supports both SQL Warehouse and All-purpose cluster types
- *All* replication options - Full Load, Store Changes and Apply Changes - loading directly into Databricks Delta tables

The new **Databricks Lakehouse (Delta)** endpoint effectively replaces the existing Microsoft Azure Databricks Delta endpoint, which will cease to be available from Replicate November 2022. Customers using the existing Microsoft Azure Databricks Delta do not need to take any actions as the new endpoint will automatically replace the existing endpoint, after upgrading to Replicate November 2022.

Amazon S3 target endpoint - AWS PrivateLink support

Customers using the **Amazon S3** target endpoint can use AWS PrivateLink to connect to EC2 instances in a virtual private cloud (VPC). To enable this functionality, a **Use AWS PrivateLink** option and **VPC Endpoint URL** field have been added to the **General** tab of the Amazon S3 target endpoint settings.

SAP HANA source endpoint - SSL support

An option to connect using SSL was added to the **General** tab of the SAP HANA source endpoint settings.

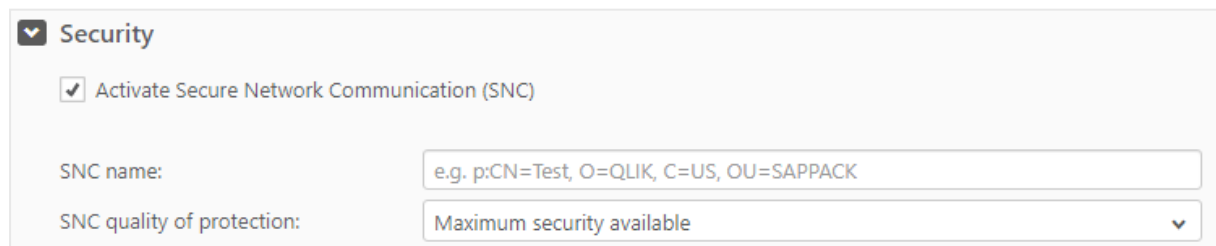
SAP HANA target endpoint - Multi tenant support

When replicating to SAP HANA target endpoint, customers can now choose **Single tenant** or **Multi tenant** according to their database architecture.

SAP Application source endpoint - SNC support

This version introduces support for connecting to the SAP Server using Secure Network Communication (SNC).

*SNC option in the SAP Application endpoint's **General** tab*



Security

Activate Secure Network Communication (SNC)

SNC name:

SNC quality of protection:

Parquet format support

Support for creating files in Parquet format has been added to the following endpoints:

- Google Cloud Storage
- Microsoft Azure ADLS

To learn about the benefits of Parquet format, see <https://databricks.com/glossary/what-is-parquet>

FIPS- compliant endpoints

FIPS compliance in Replicate endpoints depends on whether the database or ODBC vendor supports the FIPS cryptographic standard, which is not always the case. To date, the endpoints listed in the table below have been certified as FIPS-compliant. To work with these endpoints in FIPS mode, in addition to the endpoint-specific prerequisites described below, you also need to fulfill the Replicate FIPS prerequisites described in *FIPS Compliance (page 13)*.

FIPS compliant endpoints

Endpoint	Comments
File source and target	-
IBM DB2 for z/OS source	<p>The SSL connection is FIPS compliant. To connect using SSL, set the following internal parameters in the endpoint's Advanced tab:</p> <ul style="list-style-type: none"> • <code>UseSSL</code> Enables SSL • <code>SSLClientKeystoredb</code> The SSL keystore database file that is used for an SSL connection with or without the CERTIFICATE authentication. • <code>SSLClientKeystash</code> The fully qualified path of the stash file (<code>.sth</code>), which stores an encrypted password to the key database file. <p>-OR-</p> <p><code>SSLClientKeystoreDbPassword</code> The password of the keystore in string format. This password is defined when the keystore is generated.</p>
IBM DB2 for iSeries source	Supported with Replicate on Windows only. The SSL connection is FIPS compliant. To connect using SSL, set the <code>useSSL</code> internal parameter in the endpoint's Advanced tab.

Endpoint	Comments
IBM DB2 for LUW source	<p>The SSL connection is FIPS compliant. To connect using SSL, set the following internal parameters in the endpoint's Advanced tab:</p> <ul style="list-style-type: none"> • UseSSL Enables SSL • SSLClientKeystoredb The SSL keystore database file that is used for an SSL connection with or without the CERTIFICATE authentication. • SSLClientKeystash The fully qualified path of the stash file (.sth), which stores an encrypted password to the key database file. <p>-OR-</p> <p>SSLClientKeystoreDbPassword The password of the keystore in string format. This password is defined when the keystore is generated.</p>
Microsoft Azure ADLS target	Supported when using Azure Data Lake Storage (ADLS) Gen1 only.
Microsoft Azure Event Hubs target	-
Oracle source and target	For instructions on how to enable FIPS, see https://docs.oracle.com/cd/E28271_01/network.1111/e10746/asoappe.htm
Sybase ASE source	To enable FIPS, set the EnableFIPS internal parameter in the endpoint's Advanced tab to true.
Vertica target	Supported with Replicate on Linux only. Requires a FIPS compliant driver, which can be downloaded from: https://www.vertica.com/download/vertica/client-drivers/

Snowflake-based endpoints

To better reflect the expected value, the **Snowflake URL** field in the endpoint settings has been changed to **Snowflake Account/Host**.

2.2 What's new on the server side?

This section describes the new and enhanced server features.

Features and enhancements first introduced in Replicate November 2022 IR (Initial Release)

Replicate November 2022 SR1 does not include any new or enhanced engine features. This section describes the new and enhanced engine features from Replicate May 2022 Service Release 2 until and including Replicate November 2022 (Initial Release).

Checking for NOT NULL violation constraints

A new NOT NULL constraint violation handling option has been added to the **Error Handling|Data Errors** tab:

The screenshot shows the 'Data Errors Handling Policy' configuration page. The left sidebar has 'Data Errors' selected. The main content area is titled 'Data Errors Handling Policy' and includes a 'Change to Global Policy' button. Under 'For data truncation errors', there is a dropdown for 'Log record to the exceptions table'. The 'For NULL constraint violations' section is highlighted with a red box and contains three radio buttons: 'Endpoint-determined' (selected), 'Check', and 'Don't check', along with a 'Suspend table' dropdown. Below this, 'For other data errors' has a 'Log record to the exceptions table' dropdown, an 'Escalate handling when data errors reach (per table)' checkbox, a numeric input set to '0', and an 'Escalation action' dropdown set to 'Suspend table'.

For source tables, Replicate can explicitly check for NULL values in each data column of each record, and whether the same column in the corresponding target table is marked as NOT NULL. In such cases, applying the record to the target will trigger a NOT NULL constraint violation that will be difficult to recover from if the task is running in Batch Optimized Apply mode. Therefore, the record is not applied to the target. Instead, the record is handled as an error (see below) without affecting other records in the batch being applied. While this way of handling NOT NULL constraint violation is very effective, it can impact performance, and if the target database does not actually enforce NOT NULL constraint violations, Replicate can be configured to skip that check in order to save CPU time and improve performance. Because Replicate generally knows what target databases enforce NOT NULL constraint violations, it is safe to always let Replicate activate this protection based on the target endpoint type. Still, Replicate gives you the option to override this default and either force a check (with its slight overhead) or disable the check and risk NOT NULL constraint violations.



Supported with the following targets only: Google Cloud BigQuery, Amazon Redshift, Snowflake, and Microsoft Azure Synapse Analytics

DELETE and INSERT when updating a primary key column

To ease downstream processing, A new **DELETE and INSERT when updating a primary key column** has been added to the task settings' **Change Processing Tuning** tab. For tasks configured with streaming target endpoints (for example, Kafka), the DELETE + INSERT statements will be executed on the actual target tables.

For all other target endpoints, turning on **DELETE and INSERT when updating a primary key column** only impacts the Change Tables, which are created when the Store Changes replication option is enabled. Note that turning on **DELETE and INSERT when updating a primary key column** will not impact how Replicate applies Primary Key updates when the Apply Changes replication option is enabled.

Features and enhancements first introduced in Replicate May 2022 SR2

The following section describes features and enhancements first introduced in Replicate May 2022 SR2.

FIPS Compliance



This feature was first introduced in Replicate May 2022 Service Release 2.

FIPS (Federal Information Processing Standards) is a set of standards developed by the United States Federal Government for use in computer systems. FIPS 140-2 is the subset of standards which defines approved encryption algorithms used for handling sensitive information. Replicate will work in FIPS mode if the following prerequisites have been met:

- A FIPS-compliant Replicate version must be installed. To obtain the FIPS-compliant version, contact [Qlik Support](#).



Upgrading from a non FIPS-compliant Replicate is not supported. You need to uninstall your current version before installing the FIPS-compliant version. Once you have installed the FIPS-compliant version, the next Replicate version that you upgrade to does not need to be FIPS-compliant (as you will already have all the FIPS-compliant software on your machine).

- The operating system on which Replicate is installed is running in FIPS mode.

See also: [FIPS- compliant endpoints](#)

3 End of life/support features

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

3.1 End of support

Endpoint versions

Support for the following source endpoint versions has been discontinued:

- MongoDB 4.0 to 4.1
- DB2 (for IBM DB2 for z/OS) 11
- PostgreSQL 9.6
- Microsoft SQL Server 2012
- MariaDB 10.2

Support for the following target endpoint versions has been discontinued:

- PostgreSQL 9.6
- Microsoft SQL Server 2012
- MariaDB 10.2
- Microsoft Azure HDInsight 3.6
- Kafka Confluent Schema Registry 5.3.1 to 5.3.9
- Hadoop Cloudera 5.x and 6.x
- Hadoop Hortonworks 2.x
- Hortonworks Data Platform (HDP) 2.x

Endpoints

The following endpoints are no longer supported:

- Hadoop Hortonworks source
- Qlik Cloud Landing target
The functionality provided by the Qlik Cloud Landing endpoint has been superseded by Data Integration Services. For more information on Data Integration Services, see https://help.qlik.com/en-US/cloud-services/Subsystems/Hub/Content/Sense_Hub/DataIntegration/Introduction/Data-services.htm

Platforms

From Replicate May 2023, support for installing Replicate on Red Hat 7.x will be discontinued. The minimum supported version will be Red Hat 8.x.

"Backup folder preprocessing" command in the Microsoft SQL Server source endpoint

The **Backup folder preprocessing command** option has been removed from the Microsoft SQL Server source endpoint. From this version, it can only be accessed using an internal parameter.

4 Newly supported versions and third-party software

4.1 Source endpoint versions supported from Replicate November 2022 SR1

- SAP ODP version 1.1

Source endpoint versions supported from Replicate November 2022 Initial Release

The following source endpoint versions are now supported:

- MongoDB 4.2 to 5.0 and 5.x
- PostgreSQL 14
- DB2 (for IBM DB2 for z/OS) 12.1
- IBM DB2 for z/OS 2.5
- IMS 15 minor versions (for example 15.2)

4.2 Target endpoint versions supported from Replicate November 2022 SR1

- Databricks (Cloud Storage) and Databricks Lakehouse (Delta) 11.3 LTS

Target endpoint versions supported from Replicate November 2022 Initial Release

The following target endpoint versions are now supported:

- Microsoft Azure HDInsight 5.x
- Kafka - Confluent Schema Registry 5.x (starting from 5.4), 6.x and 7.x
- PostgreSQL 14
- Databricks Lakehouse (Delta) 10.4 LTS
- Confluent Cloud is now version agnostic

4.3 Target endpoint versions supported from May 2022 SR 2.

- Vertica 11.x

4.4 Windows Server 2022 support

This version introduces support for running Qlik Replicate on Windows Server 2022.

4 Newly supported versions and third-party software



First introduced in Replicate May 2022 Service Release 2.

5 Resolved issues

This section lists the resolved for the Replicate November 2022 initial release and subsequent service releases.

5.1 Resolved issues in Replicate November 2022 SR1

The following issues were resolved in Replicate November 2022 SR1.

Jira issue: RECOB-6366

Salesforce case: 60243

Type: Issue

Component/Process: Log Stream

Description: The options to automatically purge Log Stream files by time and/or size would not work properly.

Jira issue: RECOB-6327

Salesforce case: 58917

Type: Issue

Component/Process: Google Cloud BigQuery Target

Description: The Full Load task would sometimes fail with a "sqlite3: database is locked" error and no retry would occur.

Jira issue: RECOB-6381

Salesforce case: 55973

Type: Issue

Component/Process: Oracle Source - Replicate Log Reader

Description: DIRECT INSERT redo events (19.1, 11.22) would fail when using Oracle Big Endian platforms.

Jira issue: RECOB-6375

Salesforce case: 59853

Type: Issue

Component/Process: Oracle Source - Replicate Log Reader

Description: When using older Oracle versions (11.2.0.4, 12.1), redo log events 11.22 would cause unpredictable behavior.

Jira issue: RECOB-6250

Salesforce case: 58968

Type: Issue

Component/Process: Microsoft SQL Server (MS-CDC) Source

Description: When a table had no Primary Key or Index, LOB columns would be removed from the table during replication.

Jira issue: RECOB-6349

Salesforce case: 58186

Type: Issue

Component/Process: SAP Application (DB)

Description: When SAP uses SAP HANA as a backend database, columns with an empty COMPTYPE value in DD03L would not be replicated.

Jira issue: RECOB-6337

Salesforce case: 56864

Type: Issue

Component/Process: Oracle Source - Replicate Log Reader

Description: The task would sometimes report a "Unknown eye catcher in HCC DIRECT INSERT event" warning, after which some records would not be INSERTED to the target.

Jira issue: RECOB-6325

Salesforce case: 58599

Type: Issue

Component/Process: SAP Application (DB)

Description: The task would fail when a transparent table contained a column exposed as a LOB and the **Replicate LOB columns** task option was disabled.

Jira issue: RECOB-6291

Salesforce case: 59594

Type: Issue

Component/Process: SAP Application (DB)

Description: The default value of a short decimal column (0) in a Pool table would result in incorrect parsing of all subsequent columns.

Jira issue: RECOB-6287

Salesforce case: N/A

Type: Issue

Component/Process: Target endpoints: Snowflake on Amazon, Snowflake on Azure, Snowflake on Google

Description: Snowflake would not report an error when user access was disabled while the task was running.

Jira issue: RECOB-6243

Salesforce case: 54184

Type: Issue

Component/Process: SAP HANA target

Description: Excessive memory would be consumed on SAP HANA server when applying changes to tables with binary key columns.

Jira issue: RECOB-6301

Salesforce case: 56589

Type: Issue

Component/Process: Oracle Source - Replicate Log Reader

Description: The access to v\$containers permission needs to be granted when the Oracle environment is CDB/PDB (pluggable).

Jira issue: RECOB-6279

Salesforce case: 56864

Type: Issue

Component/Process: Oracle Source - Replicate Log Reader

Description: Capturing DIRECT INSERT into an HCC table would sometimes result in missing rows.

Jira issue: RECOB-6235

Salesforce case: 56864

Type: Issue

Component/Process: Oracle Source - Replicate Log Reader

Description: Capturing INSERT FROM SELECT into an HCC table would sometimes result in missing rows.

Jira issue: RECOB-6318

Salesforce case: 56486

Type: Issue

Component/Process: Salesforce (Incremental Load)

Description: Upgraded the Apache Commons Text to address security vulnerability CVE-2022-42889.

Jira issue: RECOB-6315

Salesforce case: N/A

Type: Issue

Component/Process: SAP HANA Source

Description: Added support for "max fetch interval size" for trigger-based CTS mode.

Jira issue: RECOB-6259

Salesforce case: 00034267

Type: Issue

Component/Process: MongoDB Source

Description: Replicate would fail to resume a task if document_id was not oid.

Jira issue: RECOB-5987

Salesforce case:

Type: Issue

Component/Process: Salesforce Source

Description: Tables with an unsupported data type would sometimes cause the task to restart.

Jira issue: RECOB-6239

Salesforce case: 54372

Type: Issue

Component/Process: SAP Extractor Source

Description: Changes would sometimes not be captured, with no errors or warnings in the task log.

Jira issue: RECOB-6151

Salesforce case: 51369

Type: Issue

Component/Process: SAP HANA Source

Description: When replicating from SAP HANA database version 2.0 SPS 06, captured changes would not be applied to the target.

Jira issue: RECOB-6244

Salesforce case: 50111

Type: Enhancement

Component/Process: Sybase ASE Source

Description: An option to increase the execution timeout for querying table metadata was added via an internal parameter.

Jira issue: RECOB-5983

Salesforce case: 52867

Type: Issue

Component/Process: Salesforce (Incremental Load) Source

Description: The Log Stream task would not work when the Replicate task was configured with a Salesforce Source set to Bulk API.

Jira issue: RECOB-6181

Salesforce case: 29650

Type: Issue

Component/Process: SAP Extractor

Description: The task would invoke an extractor job for a SAP table that was not a part of the task's table list.

Jira issue: RECOB-6152

Salesforce case: 12867

Type: Issue

Component/Process: Databricks on Azure Target

Description: Logging was added to assist troubleshooting intermittent data mismatch issues.

Jira issue: RECOB-5985

Salesforce case: 50008

Type: Issue

Component/Process: MySQL Source

Description: During Full Load, the table exclusion/inclusion filter would not work properly when using different cases on a database with the `lower_case_table_names` variable set to 0.

Jira issue: RECOB-5913

Salesforce case: 46218

Type: Issue

Component/Process: MySQL Source

Description: With some query platforms, CREATE DDL starting with a comment line would not be captured during CDC.

Jira issue: RECOB-5914

Salesforce case: 37728

Type: Issue

Component/Process: MySQL Source

Description: A new table added to the source would not be replicated in CDC until the task was stopped and resumed.

Jira issue: RECOB-6038

Salesforce case: 50753, 51685, 50311

Type: Issue

Component/Process: IBM for DB2 LUW Source

Description: When capturing changes from a table that was altered to be compressed (before REORG), Replicate would fail to apply the update to the target.

Jira issue: RECOB-6222

Salesforce case: 58091

Type: Issue

Component/Process: SAP HANA Source

Description: When a SAP HANA table contained active snapshots, a table partition would be loaded several times when using the "By Partitions" Parallel Load method.

Jira issue: RECOB-6526

Salesforce case: 66196

Type: Issue

Component/Process: MySQL Source

Description: It would not be possible to resume a task if MySQL Server was on Windows.

Jira issue: RECOB-6498

Salesforce case: N/A

Type: Issue

Component/Process: Logging

Description: The endpoint server (java) log would not contain any logging.

Jira issue: RECOB-6491

Salesforce case: N/A

Type: Issue

Component/Process: MongoDB Source

Description: Presentation of a 64-bit integer ("long") and a Regular Expression ("regex") has changed.

Jira issue: RECOB-6462

Salesforce case: N/A

Type: Enhancement

Component/Process: Databricks Lakehouse (Delta) Target

Description: An ADLS proxy option was added to Databricks Lakehouse (Delta). The option is enabled using an internal parameter.

Jira issue: RECOB-6499

Salesforce case: 63985

Type: Issue

Component/Process: MySQL Source

Description: Resuming a task from a CTI event, would sometimes result in missing events or/and a redundant warning message .

Jira issue: RECOB-6452

Salesforce case: N/A

Type: Issue

Component/Process: Sorter

Description: When a captured source log contained parallel processes in the same transaction (which might happen when using the SAP HANA log-based endpoint), resuming the task would sometimes not work properly. The issue was resolved using a feature flag.

Jira issue: RECOB-6500

Salesforce case: 63209

Type: Issue

Component/Process: Databricks (Cloud Storage) Target

Description: When switching from the **Access Directly** mount path option to the **Access through DBFS Mount** mount path option, the CT table in Databricks would remain empty.

Jira issue: RECOB-6484

Salesforce case: 64297

Type: Issue

Component/Process: IBM DB2 for z/OS Source

Description: If the R4Z component was paused, the task would stop and return a recoverable error when restarted or resumed.

Jira issue: RECOB-6458

Salesforce case: N/A

Type: Enhancement

Component/Process: Amazon Redshift Target

Description: Added support for Amazon Redshift PrivateLink using an internal parameter.

Jira issue: RECOB-6322

Salesforce case: N/A

Type: Enhancement

Component/Process: Databricks (Cloud Storage) and Databricks Lakehouse (Delta) Targets

Description: Added support for Databricks Unity Catalog via an internal parameter.

Jira issue: RECOB-6467

Salesforce case: 63356

Type: Issue

Component/Process: SAP Extractor Source

Description: Extractor filters would not take effect when multiple filters were defined for a single extractor.

Jira issue: RECOB-6423

Salesforce case: 56510

Type: Issue

Component/Process: Microsoft SQL Server Source

Description: When **Prevent truncation of unread changes from TLOG** was set to **Exclusively use sp_repldone within a single task** and the **Change processing mode** was set to **Online Logs Only**, BACKUP commands would sometimes fail. The issue was resolved via an internal parameter.

Jira issue: RECOB-6468

Salesforce case: 45992

Type: Issue

Component/Process: SAP Extractor Source

Description: After resuming a task, if a new extractor was activated in SAP while the task was stopped, some task extractors would not be invoked.

Jira issue: RECOB-6045

Salesforce case: 53889

Type: Enhancement

Component/Process: Salesforce (Incremental Load) Source

Description: Added support for replicating calculated fields using the Bulk API.

Jira issue: RECOB-6276

Salesforce case: 53889

Type: Issue

Component/Process: Salesforce (Incremental Load) source

Description: Full Load would sometimes fail when the SOURCE_UNLOAD logger was set to verbose.

Jira issue: RECOB-6294

Salesforce case: N/A

Type: Issue

Component/Process: Salesforce (Incremental Load) Source

Description: Changes would sometimes not be captured when there was a time difference between the Replicate machine and the Salesforce machine.

Jira issue: RECOB-6483

Salesforce case: 60089

Type: Issue

Component/Process: SAP HANA Trigger-Based Source

Description: When using trigger-based CDC, the task would fail in the following scenario:

- A source table had no primary key
- The DD03L table and the source table were defined in different schemas

Jira issue: RECOB-6465

Salesforce case: 64122

Type: Issue

Component/Process: Log Stream and Java-based Source Endpoints

Description: Performing a Full Load from Java-based source endpoints and using Log Stream would sometimes result in empty data.

Jira issue: RECOB-6415

Salesforce case: 56093

Type: Issue

Component/Process: MongoDB

Description: The DATE data type would be exposed as a string instead of a UTC epoch.

Jira issue: RECOB-6406

Salesforce case: 61543

Type: Issue

Component/Process: Kafka Target

Description: The bytes order in the Transaction ID message header field would be reversed. The issue was resolved using a feature flag.

Jira issue: RECOB-6373

Salesforce case: 54840

Type: Issue

Component/Process: ODBC Target

Description: DELETE and UPDATE operations on an IBM DB2 for LUW table without a primary key would cause the task to fail. The issue was resolved using an internal property.

Jira issue: RECOB-6441

Salesforce case: 63269

Type: Issue

Component/Process: PostgreSQL Source

Description: The capture process would not work with "half-hour" timezones (e.g. +2:30)

Jira issue: RECOB-6431

Salesforce case: 52868

Type: Issue

Component/Process: SAP HANA Log-Based Source

Description: Log events created by DMLs that changed multiple records would sometimes result in unpredictable behavior.

Jira issue: RECOB-6360

Salesforce case: 52001

Type: Issue

Component/Process: Microsoft Azure ADLS Target

Description: When mapping the BYTE data type to FIXED_LEN_BYTE_ARRAY in Parquet file format, Python and Scala Spark libraries would fail to read it. The issue was resolved with an internal parameter.

Jira issue: RECOB-6320

Salesforce case: N/A

Type: Enhancement

Component/Process: Google Cloud BigQuery Target, Databricks Lakehouse (Delta) Target, Microsoft Azure Databricks Target

Description: Added an option for Qlik Cloud Landing to configure a default proxy in the Replicate server via an internal parameter.

5.2 Resolved issues in Replicate November 2022 initial release

The following issues were resolved in Replicate November 2022 initial release:

Jira issue: RECOB-5818

Salesforce case: 44292

Type: Issue

Component/Process: SybaseIQ Target

Description: A nullable BOOLEAN data type would be created as a non-nullable BIT.

Jira issue: RECOB-6066

Salesforce case: 44742

Type: Issue

Component/Process: Engine

Description: After upgrade it would not be possible to run tasks if Replicate was installed with non-admin privileges.

Jira issue: RECOB-5804

Salesforce case: 48152

Type: Issue

Component/Process: MySQL Source

Description: Non ASCII enum values would not be captured during CDC.

Jira issue: RECOB-5893

Salesforce case: 41836

Type: Issue

Component/Process: Oracle Source - Replicate Log Reader

Description: INSERT operations would sometimes not be captured from Oracle 11.22 Redo Log event.

Jira issue: RECOB-5601

Salesforce case: 42073

Type: Issue

Component/Process: Logging

Description: With some of the cloud-based endpoints, "not found" would be written as an error message instead of a trace.

Jira issue: RECOB-5782

Salesforce case: 46218

Type: Issue

Component/Process: MySQL Source

Description: With some query platforms, CREATE DDL starting with a comment line would not be captured during CDC.

Jira issue: RECOB-5609

Salesforce case: 37728

Type: Issue

Component/Process: MySQL Source

Description: DDL changes in special formats would not be captured during CDC.

Jira issue: RECOB-5937

Salesforce case: 49312

Type: Issue

Component/Process: Oracle Source - Replicate Log Reader

Description: Incorrect INSERT operations would sometimes not be captured from Oracle 11.22 Redo Log.

Jira issue: RECOB-5689

Salesforce case: 27920

Type: Issue

Component/Process: Microsoft SQL Server Source

Description: Long VARCHAR and VARCHAR(max) values would sometimes not be replicated correctly.

Jira issue: RECOB-5877

Salesforce case: 49824

Type: Issue

Component/Process: Oracle Source - Replicate Log Reader

Description: In rare scenarios, a DELETE operation would be reported as an UPDATE by Replicate Log Reader. This would cause the target UPDATE statement to contain an empty SET CLAUSE.

Jira issue: RECOB-5763

Salesforce case: 42102

Type: Issue

Component/Process: Vertica Target

Description: The task would fail when inserting a duplicate Primary Key, even if the Apply Conflicts task policy **Duplicate key when applying INSERT** was set to **Ignore**.

Jira issue: RECOB-5861

Salesforce case: 50352

Type: Issue

Component/Process: Snowflake on AWS with Amazon S3 Staging

Description: The Full Load and CDC task would fail if the target folder had a leading slash.

Jira issue: RECOB-5799

Salesforce case: 45290

Type: Issue

Component/Process: Sorter - Kafka Target

Description: When replicating to Kafka target and the **After Full Load completes, stop the task: Before cached changes have been applied** task setting was enabled, the cached changes would be applied before the task stopped.

Jira issue: RECOB-5713

Salesforce case: 27920

Type: Issue

Component/Process: Microsoft SQL Server Source

Description: When a wide table contained long VARCHAR columns, in rare scenarios, the VARCHAR column values would be replicated to the wrong record.

Jira issue: RECOB-5810

Salesforce case: 46794

Type: Issue

Component/Process: Google Cloud BigQuery Target

Description: Empty table folders would remain in the **data_files** folder after Full Load completed.

Jira issue: RECOB-5851

Salesforce case: 41517

Type: Issue

Component/Process: IBM DB2 for z/OS Source

Description: No warning would be issued if the DATA CAPTURE CHANGES attribute was not defined for SYSIBM.SYSTABLES.

Jira issue: RECOB-5832

Salesforce case: 42768

Type: Issue

Component/Process: SAP HANA Source - Log Based Capture

Description: Excessive memory consumption would sometimes occur during the task.

Jira issue: RECOB-5842

Salesforce case: 42804

Type: Issue

Component/Process: Kafka Target

Description: In rare scenarios, the task would end abnormally.

Jira issue: RECOB-5813

Salesforce case: N/A

Type: Internal

Component/Process: Oracle Source - Replicate Log Reader

Description: The task would crash when source_log_dump logger was set to Verbose.

Jira issue: RECOB-5817

Salesforce case: 33656

Type: Issue

Component/Process: SAP HANA Source - Trigger Based Capture

Description: Changes from tables with an OID larger than 2^{32} would not be captured,

Jira issue: RECOB-5682

Salesforce case: 44211

Type: Issue

Component/Process: Google Cloud BigQuery Target

Description: When the **Create tables clustered by primary key option** was enabled in the endpoint settings, the task would fail if one of the primary key columns was a TIME data type,

Jira issue: RECOB-5830

Salesforce case: 48928

Type: Issue

Component/Process: SAP Application (DB)

Description: Change capture from clustered tables (e.g. KONV, BSEG, etc.), would sometimes fail with the following error:

```
Incorrect PART_NO 1 in the Cluster '<cluster table name>' object part 0,  
possible changes are ignored
```

Jira issue: RECOB-5820

Salesforce case: 36056

Type: Issue

Component/Process: IBM DB2 for z/OS Source

Description: IBM DB2 for z/OS would not report the lack of MONITOR2 permissions and would not capture changes.

Jira issue: RECOB-5788, RECOB-5754

Salesforce case: 41538, 43109

Type: Issue

Component/Process: Oracle Source - Replicate Log Reader

Description: The task would end abnormally during parsing of an 11.22 Redo Log event.

Jira issue: RECOB-5723

Salesforce case: N/A

Type: Issue

Component/Process: Control Tables

Description: DDLs other than ALTER TABLE would not be inserted into the DDL History Control Table (attrep_ddl_history).

Jira issue: RECOB-5768

Salesforce case: 46705

Type: Enhancement

Component/Process: Microsoft Azure Databricks Target

Description: When using a proxy connection, an option was added to disable the proxy for ADLS Gen2 storage.

Jira issue: RECOB-5752

Salesforce case: 46794

Type: Issue

Component/Process: Google Cloud BigQuery Target

Description: When using Full Load, ZIP files would sometimes be left in the **data_file** folder.

Jira issue: RECOB-5725

Salesforce case: 37794

Type: Enhancement

Component/Process: SAP HANA Source - Trigger Based

Description: Improved Left Outer Join performance when a TIMESTAMP column is used as a Primary Key.

Jira issue: RECOB-5469

Salesforce case: 38416

Type: Enhancement

Component/Process: Salesforce Source

Description: The subscribe timeout is now configurable using an internal parameter.

Jira issue: RECOB-5561

Salesforce case: 35025

Type: Issue

Component/Process: Amazon Redshift Target

Description: When working in Batch Optimized Apply mode, the task would switch to one-by-one mode instead of retrying when the following error occurred:

This type of IN/NOT IN query is not supported yet

Jira issue: RECOB-5701

Salesforce case: 44292

Type: Issue

Component/Process: CSV Target Endpoints

Description: When loading into a CSV target failed, the Full Load process would sometimes be reported as successful.

Jira issue: RECOB-5660

Salesforce case: N/A

Type: Enhancement

Component/Process: SAP HANA Trigger Based

Description: Added support for SAP HANA tables with no Primary Key, that are used in SAP with the corresponding SAP Primary Key.

Jira issue: RECOB-5680

Salesforce case: N/A

Type: Enhancement

Component/Process: Feature flags with Java SDK and auto generated endpoints

Description: Added partial support for feature flags with Java SDK and auto generated endpoints.

Jira issue: RECOB-5678

Salesforce case: N/A

Type: Enhancement

Component/Process: IBM DB2 for iSeries Source, IBM DB2 for z/OS Source, and IBM DB2 for LUW Source

Description: Added support for SSL connectivity using internal parameters.

Jira issue: RECOB-5659

Salesforce case: N/A

Type: Enhancement

Component/Process: SAP HANA Source

Description: Added support for connecting to a specific tenant database by providing the database name in the endpoint connection settings.

Jira issue: RECOB-5668

Salesforce case: N/A

Type: Enhancement

Component/Process: Sybase ASE Source

Description: Added an internal parameter for enabling FIPS support.

Jira issue: RECOB-5491

Salesforce case: 36583

Type: Issue

Component/Process: PostgreSQL Source

Description: When a timestamp record had '-infinity' as value, CDC would fail.

Jira issue: RECOB-5544

Salesforce case: 38075

Type: Issue

Component/Process: Teradata Target

Description: The TPT Log table name was being built incorrectly.

Jira issue: RECOB-5588

Salesforce case: 29650

Type: Issue

Component/Process: SAP Extractor Source

Description: The task would invoke an extractor job for a SAP table that was not a part of the task's table list.

Jira issue: RECOB-5623

Salesforce case: 38632

Type: Issue

Component/Process: SAP Hana Log-Based Source

Description: The sorter transaction swap files would accumulate and not be deleted.

Jira issue: RECOB-5568

Salesforce case: 40156

Type: Issue

Component/Process: IBM DB2 for z/OS Source

Description: This is a debug patch for the crash caused by rapid DROP/CREATE table processing.

Jira issue: RECOB-5529

Salesforce case: 34820

Type: Issue

Component/Process: SAP HANA Source Log based

Description: An empty string for both VARCHAR and NVARCHAR columns would sometimes be replicated incorrectly. Additionally, a table would sometimes be replicated incorrectly if its columns were modified.

Jira issue: RECOB-5593

Salesforce case: 41792

Type: Issue

Component/Process: Oracle Source

Description: A source transaction would fail in the Sorter when it contained changes of a captured table and SYS.OBJ\$, and was then rolled back.

Jira issue: RECOB-5599

Salesforce case: 38960

Type: Issue

Component/Process: Microsoft SQL Server Source

Description: When "Read the backup log directly from a file" was enabled, the backup log usage would always be prioritized, even if the online logs priority was configured.

Jira issue: RECOB-5621

Salesforce case: 33656

Type: Issue

Component/Process: Metadata Manager

Description: Events from some tables would not be captured when the table ID of the number of captured tables was greater than 2*32.

Jira issue: RECOB-5459

Salesforce case: N/A

Type: Enhancement

Component/Process: SAP Application (DB) Source

Description: Business Objects changes processing was improved.

Jira issue: RECOB-5330

Salesforce case: 12867

Type: Issue

Component/Process: Microsoft Azure Databricks Target

Description: Added logging to troubleshoot a count difference between source and target issue.

Jira issue: RECOB-5307

Salesforce case: 34640

Type: Issue

Component/Process: Sybase ASE source

Description: The FLOAT data type would be replicated incorrectly during both Full Load and CDC.

Jira issue: RECOB-5487

Salesforce case: 00036960

Type: Issue

Component/Process: Salesforce Source

Description: When only the LastModifiedDate field was changed, the change would be ignored.

Jira issue: RECOB-5212

Salesforce case: 24055

Type: Issue

Component/Process: PostgreSQL Source

Description: Amazon RDS for PostgreSQL would not handle network issues properly.

Jira issue: RECOB-5522

Salesforce case: N/A

Type: Issue

Component/Process: Replicate General

Description: Replicate on Linux was built using FIPS mode

Jira issue: RECOB-5527

Salesforce case: 34091

Type: Issue

Component/Process: IBM DB2 for z/OS Source

Description: When a captured table had no CAPTURE DATA CHANGES attribute, the corresponding warning messages would not be issued.

Jira issue: RECOB-5484

Salesforce case: 3898

Type: Issue

Component/Process: Google Cloud Storage

Description: When a task was stopped and then resumed, the CDC files generated before the task was stopped would not be uploaded to Google Cloud.

Jira issue: RECOB-5511

Salesforce case: **Jira issue:** RECOB-5609

Salesforce case: 37728

Type: Issue

Component/Process: MySQL Source

Description: DDL changes in special formats would not be captured during CDC.

Type: Issue

Component/Process: MySQL (MariaDB)

Description: A new table added to the source would not be replicated in CDC until the task was stopped and resumed.

Jira issue: RECOB-5564

Salesforce case: 38064

Type: Issue

Component/Process: Salesforce Source (Incremental Load)

Description: Timeout errors would sometimes occur during Full Load.

Jira issue: RECOB-5237

Salesforce case: N/A

Type: Issue

Component/Process: Google BigQuery Target

Description: When **Create tables clustered by primary key** was enabled in the endpoint's **Advanced** tab, and **Apply changes using SQL MERGE** was turned on in the task settings, INSERTs would expose the data in the log.

Jira issue: RECOB-5509

Salesforce case: 21512

Type: Enhancement

Component/Process: Log Stream

Description: When using Log Stream and the **Source change position (e.g. SCN or LSN)** field is set to **0**, the task will resume from the first event of the latest timeline.

Jira issue: RECOB-5176

Salesforce case: 30284

Type: Issue

Component/Process: Microsoft Azure SQL (MS-CDC) source

Description: When the date format in the default language was not MDY, Full Load would fail.

Jira issue: RECOB-5485

Salesforce case: N/A

Type: Issue

Component/Process: Amazon S3 Target

Description: The **Region code** UI label was incorrectly named **Bucket region**.

Jira issue: RECOB-5364

Salesforce case: 33794

Type: Issue

Component/Process: Snowflake Azure Target

Description: When the **Optimize Inserts** option was enabled in Batch Optimized Apply mode and JSONandXMLColumns-SnowflakeTarget was turned on (the default), the task would fail when trying to insert a CLOB column whose length exceeded 65535 characters.

Jira issue: RECOB-5445

Salesforce case: 21090

Type: Issue

Component/Process: Salesforce Source

Description: The task was stop capturing changes every few weeks, with the following error (excerpt)

exception=java.lang.OutOfMemoryError: Java heap space

Jira issue: RECOB-5457

Salesforce case: 34455

Type: Issue

Component/Process: SAP HANA Source

Description: SAP HANA views would not be replicated after upgrading.

Jira issue: RECOB-5540

Salesforce case: 38064

Type: Issue

Component/Process: Salesforce Source (Incremental Load)

Description: A timeout error would sometimes occur when unloading a table.

Jira issue: RECOB-5532

Salesforce case: 34641

Type: Issue

Component/Process: SAP HANA Source

Description: Filtering archiving user names that contained a period would not work.

Jira issue: RECOB-5232

Salesforce case: 4844

Type: Issue

Component/Process: Oracle Source - Replicate Log Reader

Description: Replicate Log Reader would crash after parsing a Direct Insert Redo Log Event 19.1.

Jira issue: RECOB-5517

Salesforce case: 37119

Type: Issue

Component/Process: Oracle Source - Replicate Log Reader

Description: INSERT operations in a compressed table would be missing when Goldengate was enabled in Oracle.

Jira issue: RECOB-5072

Salesforce case: NA

Type: Issue

Component/Process: Google BigQuery Target, Security

Description: When **Create tables clustered by primary key** was enabled in the endpoint's **Advanced** tab, and **Apply changes using SQL MERGE** was turned on in the task settings, INSERTs would expose the Primary Key data in the log.

Jira issue: RECOB-5180

Salesforce case: 31399

Type: Issue

Component/Process: SAP Application (DB) Source with a Microsoft SQL Server backend

Description: When trying to capture changes from a Cluster table, the task would sometimes crash with the following error:

```
Incorrect PART_NO <> in the Cluster '<>' object part <>, possible changes are ignored
```

Jira issue: RECOB-5439

Salesforce case:

Type: Issue

Component/Process: Salesforce Source (Streaming CDC)

Description: Parsing of a custom time field would fail during Full Load with the following error:

```
Unable to parse '<time>' as TIME. Will use null instead.
```

Jira issue: RECOB-5394

Salesforce case:

Type: Issue

Component/Process: Salesforce Source (Incremental Load)

Description: A "Read timed out" error would sometimes occur when working with Bulk API.

Jira issue: RECOB-5501

Salesforce case: 38566

Type: Issue

Component/Process: Oracle Source - Replicate Log Reader

Description: When the following settings were enabled in the endpoint's Advanced tab, the redo files would not be deleted automatically.

- Copy redo logs to temporary folder
- Replicate has file level access to temporary folder
- Access archived redo logs in folder
- Delete processed archived redo log files

Jira issue: RECOB-5368

Salesforce case: 27920

Type: Issue

Component/Process: Microsoft SQL Server Source

Description: When a wide SQL Server table contained long VARCHAR columns, in rare scenarios, the VARCHAR column values would be replicated to the wrong record.

Jira issue: RECOB-5352

Salesforce case: 34292

Type: Issue

Component/Process: Salesforce Source (Incremental Load)

Description: Confusing warning messages would be shown in certain scenarios.

Jira issue: RECOB-5269

Salesforce case: 37942

Type: Issue

Component/Process: Salesforce Source (Incremental Load)

Description: Compound field types such as Address would replicated as NULL when using SOAP API.

Jira issue: RECOB-5321

Salesforce case: 34832

Type: Issue

Component/Process: Salesforce Source (Incremental Load)

Description: Remove column DDL cause to recoverable error loop

Jira issue: RECOB-5378

Salesforce case:

Type: Issue

Component/Process: Salesforce Source (Incremental Load)

Description: Unicode characters would not be replicated when using BULK API.

Jira issue: RECOB-5324

Salesforce case: 34585

Type: Issue

Component/Process: Oracle Source - Replicate Log Reader

Description: When starting the task from a timestamp, Replicate Log Reader would sometimes capture Dynamic Partition table operations from the wrong Dynamic Partition table.

Jira issue: RECOB-5472

Salesforce case: 37942

Type: Issue

Component/Process: Salesforce Source (Incremental Load)

Description: Full Load would sometimes fails when a task was resumed.

Jira issue: RECOB-5221

Salesforce case: 32297

Type: Issue

Component/Process: MySQL Homogeneous

Description: ALTER TABLE operations would result in an unexpected character set being defined on the target.

Jira issue: RECOB-4881

Salesforce case: 02101374

Type: Issue

Component/Process: Microsoft SQL Server Source (from 2016)

Description: When the Primary Key was DATETIME, LOB lookup would fail, resulting in missing LOB columns.

Jira issue: RECOB-5242

Salesforce case: 33124

Type: Issue

Component/Process: PostgreSQL Source

Description: The task would sometimes crash when a large number of tables were truncated in a single operation.

Jira issue: RECOB-4880

Salesforce case: 24364

Type: Issue

Component/Process: Teradata Source

Description: When using Parallel Load, in a task defined with the Full Load and Apply Changes options, the task would sometimes crash during Full Load.

Jira issue: RECOB-5486

Salesforce case: 37690

Type: Issue

Component/Process: Snowflake

Description: Incorrect default values were assigned to some of the internal parameters.

Jira issue: RECOB-5435

Salesforce case: 36027

Type: Issue

Component/Process: SAP Application (DB)

Description: When a captured cluster table was changed, SAP Client records that were not configured in SAP Application (DB) would be replicated.

Jira issue: RECOB-5448

Salesforce case: NA

Type: Enhancement

Component/Process: Replicate Engine

Description: Updated the Kerberos OpenSSL libraries.

Jira issue: RECOB-5276

Salesforce case: 33320

Type: Issue

Component/Process: Microsoft Databricks Azure

Description: When the Sequence storage format was used and the files were compressed with gzip, empty files would be left in the Databricks partition if the load operation was interrupted.

Jira issue: RECOB-5427

Salesforce case: 37749

Type: Issue

Component/Process: Amazon Redshift Target

Description: It would not be possible to save the endpoint settings without changing the default Bucket region.

Jira issue: RECOB-5327

Salesforce case: 32256

Type: Issue

Component/Process: SAP HANA Trigger Based Capture

Description: When using the Commit-Timestamp internal parameter, latency issues would be encountered when processing changes from non-captured tables.

Jira issue: RECOB-5419

Salesforce case: 33783

Type: Issue

Component/Process: Microsoft SQL Server Source

Description: The LSN on a secondary replica in an AlwaysOn environment could not be located.

Jira issue: RECOB-5372

Salesforce case: 34751

Type: Issue

Component/Process: Installer

Description: The Qlik Replicate installation timeout was increased to 300 seconds.

Jira issue: RECOB-5375

Salesforce case: 31588

Type: Issue

Component/Process: SAP DB

Description: After upgrading to Replicate 2021.11 and resuming a task, the first DDL change would cause the associated table to be suspended.

Jira issue: RECOB-5068

Salesforce case: 30567

Type: Issue

Component/Process: BigQuery TARGET

Description: When **Create tables clustered by primary key** was enabled in the endpoint's **Advanced** tab, and **Apply changes using SQL MERGE** was turned on in the task settings, INSERTs would fail when the Primary Key started or ended with a single quote.

Jira issue: RECOB-5170

Salesforce case: 23555

Type: Issue

Component/Process: Amazon Redshift

Description: Replicate would not attempt to reconnect to Amazon Redshift if the connection was lost.

Jira issue: RECOB-5380

Salesforce case: 25486

Type: Issue

Component/Process: Snowflake

Description: In Batch Optimized Apply mode, explicit COMMITs would be performed when applying changes to Snowflake, resulting in unacceptable costs.

Jira issue: RECOB-5164

Salesforce case: 31074

Type: Issue

Component/Process: Control Tables

Description: When the DDL History Control table was enabled, and a column was changed from TEXT to VARCHAR(max) or vice versa, the task would encounter a recoverable error.

Jira issue: RECOB-5342

Salesforce case: 30113

Type: Issue

Component/Process: SAP DB

Description: A special SAP date using the year 0101 would not be replicated correctly.

Jira issue: RECOB-5268

Salesforce case: 28583

Type: Issue

Component/Process: SAP DB

Description: Changes would sometimes not be captured when replicating from SAP cluster tables with Microsoft SQL Server as the underlying database.

Jira issue: RECOB-5174

Salesforce case: 32472

Type: Issue

Component/Process: LOB Support

Description: LOB support would be disabled when the **No record found for applying an update → INSERT the missing target record** error handling option was selected, AND a table was changed with a transformation.

6 Known Issues

This section describes the known issues in Replicate November 2022 SR (Service Release) 1 and in Replicate November 2022 IR (Initial Release).



Known issues in Replicate November 2022 IR have been resolved in Replicate November 2022 SR 1.

6.1 Known issues in Replicate November 2022 SR 1

This section lists the known issues in Replicate November 2022 Service Release 1.

Jira issue: RECOB-6617

Salesforce case: 64536

Type: Issue

Component/Process: Sorter

Description: When a task is configured with SAP Application (DB), reloading a table might cause duplicate keys if a table is reloaded while changes are being applied to it.

Jira issue: RECOB-6618

Salesforce case: 67786

Type: Issue

Component/Process: SAP HANA Source (Log Based)

Description: When the ROWID key value is absent, DELETE operations sometimes fail on the target and INSERT operations might be missed.

Jira issue: RECOB-6604

Salesforce case: 65993

Type: Issue

Component/Process: Snowflake Azure Target

Description: When specifying an Azure blob staging folder without a slash ('/'), Replicate creates the folder without its first letter.

Jira issue: RECOB-6581

Salesforce case: 66709

Type: Issue

Component/Process: SAP HANA Source (Log Based)

Description: Tables are sometimes suspended with the following warning message "Failed to get data values for ... data is set to NULL".

Jira issue: RECOB-6370

Salesforce case: 57757

Type: Issue

Component/Process: Oracle Source - Replicate Log Reader

Description: When Oracle is restarted with the RESETLOGS command, the task fails.

Jira issue: RECOB-6567

Salesforce case: 66516

Type: Issue

Component/Process: SAP HANA Source (Log Based)

Description: Tables are sometimes suspended without any explanation.

Jira issue: RECOB-6532

Salesforce case: 63213

Type: Issue

Component/Process: Databricks Lakehouse (Delta) Target

Description: When the **Apply batched changes to multiple tables concurrently** option is enabled, the task occasionally fails to apply the changes but succeeds after retries.

6.2 Known issues in November 2022 Initial Release

This section lists the known issues in Replicate November 2022 IR (Initial Release).

Jira issue: RECOB-5913

Salesforce case: 46218

Type: Issue

Component/Process: MariaDB Source

Description: A CREATE DDL change starting with comment line will not be captured.

Jira issue: RECOB-5597

Salesforce case: 37728

Type: Issue

Component/Process: MySQL Source

Description: DDL changes in special formats are sometimes not captured during CDC.

Jira issue: RECOB-6152

Salesforce case: 12867

Type: Issue

Component/Process: Microsoft Azure Databricks Target

Description: The number of replicated records does not always correlate with the number of source records.

Jira issue: RECOB-6181

Salesforce case: 29650

Type: Issue

Component/Process: SAP Extractor Source

Description: The SAP Extractor endpoint sometimes activates an extractor job that is not part of the CDC task.

Jira issue: RECOB-5985

Salesforce case: 50008

Type: Issue

Component/Process: MySQL Source

Description: Table exclusion/inclusion logic during CDC is case sensitive.

Jira issue: RECOB-5983

Salesforce case: 52867

Type: Issue

Component/Process: Salesforce (Incremental Load) Source

Description: It is not possible to work with Log Stream using the Bulk API.

Jira issue: RECOB-6038

Salesforce case: 50753, 51685, 50311

Type: Issue

Component/Process: IBM for DB2 LUW Source

Description: When capturing changes from a table that was altered to be compressed (before REORG), Replicate fails to apply the update to the target.

Jira issue: RECOB-6151

Salesforce case: 51369

Type: Issue

Component/Process: SAP HANA Source - Log Based CDC

Description: Captured changes are sometimes not replicated to the target.

Jira issue: RECOB-6222

Salesforce case: 58091

Type: Issue

Component/Process: SAP HANA Source

Description: The SAP HANA partition list contains duplicate partition IDs when a table contains active snapshots.