

Qlik Replicate 6.5 Release Notes - November 2019

Qlik Replicate 6.5 introduces several new features and enhancements, as well as exciting additions to the Support Matrix including:

- » **Salesforce** as a SaaS application source
- » **MongoDB** as a NoSQL-based source
- » **Google Cloud BigQuery** as a data warehouse target
- » **Databricks on Azure** as a cloud target that supports fully managed Spark clusters

Additional features include significant performance improvements, security hardening, support for the new generation ADLS Gen2 Azure Cloud Storage, and support for multiple endpoint configurations (useful in the event of failover or migration).

Note 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](#)
 - » [Performance Improvements](#)
 - » [Security Hardening](#)
 - » [Microsoft SQL Server Source Endpoint Enhancements](#)
 - » [Oracle Source Endpoint Enhancements](#)
 - » [Microsoft Azure SQL Managed Instance Source Endpoint Enhancement](#)
 - » [SAP HANA Source Multi-Tenant Support](#)
 - » [Support for Multiple Configurations of the Same Endpoint Type](#)
 - » [Azure Data Lake Storage \(ADLS\) Gen2 Support](#)
 - » [Other Enhancements/Changes](#)
- » [Newly Supported Endpoints](#)
- » [Newly Supported Endpoint Versions](#)
- » [End of Life/Support and Deprecated Features](#)
- » [Resolved Issues and Customer Requested Enhancements](#)

Migration and Upgrade

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

Controlling Execution of User-Defined Commands

As part of the security improvements to this version, execution of user-defined commands is blocked by default in a clean installation. When pre or post processing of data files is required, user-defined commands can be enabled.

However, in order to maintain backward compatibility when upgrading from previous versions, the `enable_execute_user_commands` parameter in the `repctl.cfg` file will be set to "true" (allowing execution of user-defined commands).

For instructions on changing this parameter, please refer to the Replicate Help.

For more information on this feature, see [Security Hardening](#).

Regenerating the Replicate Self-Signed Certificate after Upgrade

Due to security enhancements, after upgrading, customers that are using Replicate's self-signed certificate (i.e. instead of their own certificate) need to perform the following procedure in order to generate a new self-signed certificate.

1. Delete all the `*.pem` files from `<replicate_installation_folder>data/ssl/data`.
2. Restart the **Attunity Replicate Server** service.

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 6.5.
3. Upgrade the transports as described in the Replicate Help.
4. Restart the tasks.

Oracle Source Endpoint

Upgrade from Replicate 5.5

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 Attunity Replicate.

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

Replicating from a RAC Environment

Customers replicating from a RAC environment are strongly advised to upgrade to Replicate 6.5; otherwise, changes performed on the non-primary node during Full Load may not be captured. After upgrading, you also need to grant new permissions to the Replicate user.

For more information, see [Oracle Source Endpoint Enhancements](#).

IBM DB2 for z/OS Source Endpoint

Due to enhancements made to the Attunity R4Z (previously known as R4DB2) installation procedure as well as the underlying logic, upgrading the Attunity 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 *Attunity R4Z Installation and Configuration Guide*.

Upgrading the Attunity Replicate Console Only

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

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

Workaround:

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

```
RepUiCtl.exe repository upgrade --repository ..\Data\GlobalRepo.sqlite
```

4. Restart the Attunity 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 6.5 (from Replicate 5.5) should contact Attunity Support.

Log Stream Target Endpoint

Due to the new compression slider, after upgrading to Replicate 6.5, the compression level will be changed from 9 (the highest) to 2 (the lowest and the new default). This will result in larger file sizes on the disk.

For more information on this feature, see [Other Enhancements/Changes](#).

Tasks that were Started from Timestamp

Note Applies when upgrading from Replicate 5.5 or 6.3 only.

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.

Attunity Enterprise Manager (AEM) Compatibility

Replicate 6.5 is compatible with AEM 6.5 only.

New and Enhanced Features

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

Performance Improvements

- » **Kafka Target Endpoint** - Performance was improved when using Avro format to encode the messages produced by the Kafka target endpoint.
- » **Log Stream Target Endpoint** - A new compression level slider has been added to the **General** tab. You can adjust the slider according to your needs. A higher compression level will conserve disk space (at the expense of latency) whereas a lower compression level will increase processing speed.
- » **Microsoft SQL Server Source Endpoint** - Performance has been improved by reading LOBs directly from the transaction logs instead of performing a lookup.
- » **Extended Parallel Load Support** - Support for the Parallel Load feature has been extended to the following endpoints:
 - » Teradata Source
 - » Amazon RDS for Microsoft SQL Server Source
 - » Microsoft Azure Data Warehouse Target
 - » Snowflake on AWS Target
 - » Snowflake on Azure Target
- » **Change Processing** - Support for the **Apply batched changes to multiple tables concurrently** option has been extended to the following target endpoints:
 - » Amazon Redshift
 - » Microsoft Azure SQL Data Warehouse

Security Hardening

This version includes the following security enhancements:

- » Mail Settings: A **Use SSL** option was added to the mail server settings (which are used to send notifications)
 - » When **Use SSL** is enabled, user can also select the new **Verify peer** and/or **Verify host** fields to validate the server certificate.
- » On Windows, it's now possible to run Replicate under a non-privileged account. To do so, follow the instructions in the "**Running Replicate under a Different Account**" topic in the Replicate Help.

- » With new installations, the server password must now satisfy the newly implemented strong password criteria. For more information about the strong password criteria as well as instructions for generating a strong password, refer to the **"Changing the Server Password topic"** in the Replicate Help.
- » User-defined commands in endpoint settings:
 - » Several endpoints support running user-defined commands for pre/post-processing files. In the new **Operating System-Level Credentials** tab (in Server view), you can provide alternate user credentials for any user-defined command that needs to be executed at operating system level.

By default, such commands are executed under the Replicate Server service account on the Replicate Server machine. This may constitute a security risk, as it allows any Replicate user with Admin or Designer permissions to specify user-defined commands that could lead to a full compromise of the server.
 - » With new installations, the execution of user-defined commands is blocked by default. To allow the execution of such commands, follows the procedure described in the **"Allowing Execution of User-Defined Commands"** topic in the Replicate Help.
- » Support for TLS 1.0 and TLS 1.1 has been deprecated.

Microsoft SQL Server Source Endpoint Enhancements

From Replicate 6.5:

- » You can now use a non-sysadmin user when replicating from Microsoft SQL Server configured with AlwaysOn.

Note Due to changes to the required procedure for a non-sysadmin user, if you are upgrading from Replicate 6.4, you need to re-execute the non-sysadmin setup scripts. See the Replicate Help for the relevant scripts.

- » Dynamic Data Masking is now supported.

Oracle Source Endpoint Enhancements

Replicate 6.5 introduces the following changes in functionality:

- » RESETLOGS operations are now captured by default.
- » In previous versions, Replicate would detect open transactions on the RAC master (primary) node only, which would sometimes result in changes not being captured. From this version, Replicate is now able to detect open transactions on *all* RAC nodes.

Customers wishing to benefit from this new functionality need to grant the Replicate user the following permissions:

- » **Regardless of which method is used to access the redo logs:**

Grant SELECT ON GV_\$TRANSACTION (for detecting open transactions on *all* RAC node)

Note that once this permission is granted, the SELECT ON V_\$TRANSACTION permission is no longer needed.

- » **When using Attunity Log Reader to access the redo logs:**

Grant SELECT ON V_\$DATABASE_INCARNATION (for capturing RESETLOGS operations)

Note If Replicate detects that the user (specified in the endpoint settings) does not have the requisite permissions, a warning will be written to the log and the task will continue as normal. However, when replicating from a RAC environment and/or capturing RESETLOGS operations, *not setting these permissions may result in data loss and/or unpredictable behavior.*

Microsoft Azure SQL Managed Instance Source Endpoint Enhancement

You can now use a non-sysadmin user with the Microsoft Azure SQL Managed Instance source endpoint.

SAP HANA Source Multi-Tenant Support

The SAP HANA source endpoint now supports replication from both single tenant and multi-tenant architectures (on a single host).

Support for Multiple Configurations of the Same Endpoint Type

Starting from Replicate 6.5, it is now possible (using the AEM SDK) to define multiple configurations of a source or target endpoint and switch them as required. Such functionality may be useful in the event of database failover or when migrating from a testing to a production environment, for example.

For more information, please refer to the *Attunity Enterprise Manager Developer's Guide*.

Azure Data Lake Storage (ADLS) Gen2 Support

Support for Azure Data Lake Storage (ADLS) Gen2 was added to the following target endpoints:

- » Microsoft Azure SQL Data Warehouse
- » Microsoft Azure ADLS
- » Microsoft Azure Databricks
- » Microsoft Azure HDInsight

Other Enhancements/Changes

- » **Full Load Recovery** - If the connection is reset by the source or target endpoints during Full Load (e.g. due to a lengthy idle time), Replicate will now attempt to re-establish the connection while the task is still running (i.e. without stopping and restarting the task).
- » **MemSQL Target** - Support for creating columnstore tables was added. When the new **Create columnstore tables** option is selected in the **Advanced** tab, Replicate will create the target tables in columnstore format (instead of rowstore format).
- » **MySQL Source** - When there is a Unique Index on nullable segments but no Primary Key, a Unique Index will be created on the target.
- » **SAP Application and SAP Application (DB) Source Endpoints** - Replicate 6.5 introduces two new "delete" transports (**DeleteCodeCommon.zip** and **DeletePackageCommon.zip**). These new transports need to be applied if you wish to uninstall the Attunity Replicate for SAP Client, unless Attunity Gold Client is running in the same environment. For usage instructions, please refer to the Replicate Help.

Newly Supported Endpoints

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

Notes The Salesforce source, MongoDB source, and Google BigQuery target endpoints are currently in beta.

MongoDB Source Endpoint

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

Salesforce Source Endpoint

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

Google SQL Cloud for MySQL Source

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

Google SQL Cloud for PostgreSQL Source

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

Amazon RDS for MariaDB Source

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

Google Cloud BigQuery Target Endpoint

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

Microsoft Azure Databricks Target Endpoint

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

Newly Supported Endpoint Versions

The following endpoint and platform versions are now supported:

Endpoints:

- » PostgreSQL 10.7, 11.1, 11.2, and 11.3 source and target
- » Amazon RDS for PostgreSQL target 10.7 and 11.2
- » Sybase ASE source is now supported with Replicate on Linux
- » Pivotal Greenplum 5.16 and 5.7 target
- » Vertica 9.1 target
- » Kafka target 2.1, 2.2, and 2.3
- » Teradata 16.20 target
- » Cloudera 6.1
- » Oracle 19.x target

End of Life/Support and Deprecated Features

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

Attunity Replicate 6.3

Attunity's Support and Maintenance Agreement guarantees support for three Replicate versions, as follows:

- » When the current Version Release *does not* have a minor release (i.e. x.0), then ATTUNITY shall support:
 - » The current Version Release (1)
 - » The last two (2) Minor Releases of the previous Major Release
- » When the current Version Release *does* have a minor release - for example x.1 - then ATTUNITY shall support:
 - » The last two (2) Minor Releases of the current Major Release
 - » The last one (1) Minor Release of the previous Major Release

Based on the policy outlined above, Attunity announces the end of life of Replicate 6.3. Customers who need additional time to upgrade, can do so during the six-month period following the GA release of Replicate 6.5. Nevertheless, Replicate customers are encouraged to upgrade to Replicate 6.5 well before the end of the six-month period. No additional maintenance charges will be introduced during this period.

Attunity understands that in certain circumstances a customer may need an additional support extension to ensure business continuity. If this is the case, please contact Attunity Support to discuss the available options.

Deprecated Support for Endpoint Versions

Microsoft SQL Server 2005 will no longer be supported as a source or target endpoint.

Deprecated Support for Linux Platforms/Versions

From Replicate 6.5, the following platforms/versions are no longer supported:

- » SUSE Linux
- » RedHat versions 6.7 and earlier, and 7.1-7.3

Resolved Issues and Customer Requested Enhancements

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

Component/Process	Type	Description	Ref #
<ul style="list-style-type: none"> » Amazon EMR » Microsoft Azure HDInsight » Google Dataproc » Hortonworks Data Platform 	Issue	<p>Target endpoints would incorrectly map the TIME data type to TIMESTAMP instead of STRING.</p> <p>New tables will be created correctly. For existing tables, modify the column type in Hive accordingly.</p>	N/A
PostgreSQL Source to PostgreSQL Target	Issue	<p>When working with unlimited LOB support, the task would fail when loading tables with columns of the following data types:</p> <ul style="list-style-type: none"> » circle » polygon » path 	183412
Oracle Homogeneous	Issue	<p>When the Oracle source endpoint was set to: Expose NUMBER as:VARCHAR and the source column data exceeded precision 32, the task would fail during Full Load.</p>	185286

Component/Process	Type	Description	Ref #
ODBC Target		<p>When replicating to HP NonStop SQL/MP and working with an Audit table, the following error would be encountered due to incorrect handling of a non-standard TIMESTAMP format:</p> <p>Invalid scale value</p> <p>The issue was resolved using an internal parameter.</p>	185406
Microsoft SQL Server Source	Issue	<p>When the Microsoft SQL Server endpoint's Replicate has file-level access to the backup log files option was enabled, parsing of an operation in a large row would sometimes cause the task to fail.</p> <p>Now, in such cases, the task will no longer fail. Instead, the relevant table will be suspended with additional verbose logging.</p> <p>Customers will then be able to reload the suspended table to overcome the issue.</p>	186503
Microsoft SQL Server	Issue	Out of row data would not be converted to UTF-8 format.	186503
Microsoft SQL Server to Kafka Target	Issue	The BEFORE-IMAGE for tables altered with ADD COLUMN would sometimes contain corrupted data.	186503

Component/Process	Type	Description	Ref #
Kafka Target	Issue	<p>The task would fail and get stuck in an infinite loop when the input DOUBLE data type was either NAN or INFINITY.</p> <p>Now, a data error will be reported and a record will be inserted into the Apply Exceptions table.</p>	186503
Microsoft SQL Server Source	Issue	<p>When the Microsoft SQL Server endpoint's Replicate has file-level access to the backup log files option was enabled, parsing of an operation in a large row would sometimes cause the task to fail with the following error.</p> <pre>Data length is '8073' that is longer than allowed 8060. Data index is '0'</pre> <p>This issue was resolved by increasing the maximum data buffer size in the code.</p>	187370
MySQL Source	Issue	<p>The task would sometimes stop unexpectedly after failing to read the binlog.</p>	187442

Component/Process	Type	Description	Ref #
SAP Application	Issue	<p>Spaces would be removed from Primary Key values that were padded with spaces in the SAP application database.</p> <p>Spaces will now be added whenever the Primary Key value isn't long enough to match the length definition in the SAP metadata.</p>	187519
Teradata Target	Issue	<p>During Full Load, DATETIME values in the following format would not be replicated to the target:</p> <p>9999-12-31 XX:XX:XX</p>	187579
MySQL Source	Issue	<p>Multiple issues would occur when replicating a large auditing table from MySQL source in multiple environments.</p>	187821
ODBC Infrastructure	Issue	<p>Incorrect conversion of an internal string to the TIME data type would result in the following error:</p> <p>Duplicate entry '1149740' for key 'PRIMARY' during migration</p>	187826
Snowflake Target	Issue	<p>During CDC, some changes would not be applied to the target due to a connectivity issue.</p>	187966

Component/Process	Type	Description	Ref #
Error Handling		Multiple read operation errors would occur due to forcefully disconnecting the client.	188087
Oracle Source - Binary Reader	Issue	CDC tasks would stop unexpectedly when replicating from an Oracle standby environment.	188139
Sybase ASE Source	Issue	Merge operations would not be supported by the SAP Sybase ASE source endpoint while running CDC resulting in a fatal error.	188301
SAP Application (DB)	Issue	An Apply Changes only task would stop unexpectedly when replicating a DD07T (client independent) table.	188343
Microsoft SQL Server	Issue	When trying to resolve a partition ID that no longer exists, CDC would fail with the following error: Failed in resolving object ID by partition ID	188384
Hadoop Target	Issue	On rare occasions, an error would occur when using Sequence format with GZIP compression.	188415

Component/Process	Type	Description	Ref #
Oracle Source - Attunity Log Reader	Issue	<p>High CDC latency would be encountered when replicating from an Oracle 12.2 source with ASM. This was partly due to the inability to specify a retry interval less than 1 second.</p> <p>Latency was reduced by allowing "0" (approximately 10 milliseconds) to be specified in the Retry interval field in the Advanced tab.</p>	188444
SAP HANA Source	Issue	<p>Capturing changes from a table with a name exceeding 10 characters would lead to unpredictable results due to insufficient space in the internal buffer.</p>	188546
ODBC Target (AIS SQL/MP)		<p>Sending an ODBC wide string parameter which was not null terminated to the AIS ODBC driver would result in corrupted data being inserted into the <code>change_record</code> of the AUDIT table.</p>	188585
Hadoop Target	Issue	<p>Upload of the CSV file to HDFS would fail because the file size was zero.</p>	188762
REST	Issue	<p>Front end UI would be unresponsive due to memory leak and error handling problems.</p>	188773

Component/Process	Type	Description	Ref #
Oracle Source - Attunity Log Reader	Issue	During CDC, the task would sometimes stop unexpectedly when a table contained a LONG value exceeding 8000 bytes and the value was updated to an even longer value (e.g. 16,000 bytes).	188775
Repository	Issue	After upgrading to Replicate 6.3, resuming a task with transformations would reload the tables.	188862
PostgreSQL Source	Issue	Incorrect implementation of TRACE logging would sometimes result in extremely large logs being created.	188935

Component/Process	Type	Description	Ref #
SAP HANA Source	Issue	<p>An incremental value (indicating a change) would be reserved for the INDX column in the attrep_cdc_changes table before the actual transaction changes were committed to the source database.</p> <p>This would sometimes result (usually with very long transactions) in some of the source changes not being applied to the target, as Replicate would continue reading subsequent changes from the attrep_cdc_changes table without waiting for the previous transaction to complete in the source database.</p> <p>An internal parameter was added that allows customers to set the anticipated transaction length as desired (default 60 seconds).</p> <p>Logging was also enhanced to report missing changes or rolled back transactions.</p>	188936
Repository	Issue	<p>Replicate would fail to start after upgrading from version 6.1 to version 6.3.</p>	188941

Component/Process	Type	Description	Ref #
Oracle Source -Attunity Log Reader	Issue	The task would fail after the following warning: Redo log contains both encrypted and unencrypted events for table <table_name>	189074
Oracle to Oracle	Issue	When VARCHAR columns contained numeric values in the Oracle source, and the data type of the corresponding Oracle target column was NUMBER, not all valid NUMBER values would be migrated correctly.	189125
File Channel	Issue	When the remote File Channel task was set to "Store Changes" only, columns added as the result of an "Add Column" transformation (defined in the remote task) would not be populated with data if an "ALTER TABLE ADD COLUMN" operation occurred in the source.	189180

Component/Process	Type	Description	Ref #
File Channel	Issue	<p>When set to "Store Changes" only, the remote File Channel task would fail under the following circumstances:</p> <ul style="list-style-type: none"> » A DELETE operation was performed on a record with a LOB column » An "Add Column" transformation was defined (in the remote task). 	189224
ODBC Target	Issue	<p>Replicating tables without a primary key to DB2z via the ODBC target endpoint would cause the error:</p> <pre>SQL0417N A statement string to be prepared contains parameter markers as the operands of the same operator</pre>	189460
Hadoop Target	Issue	<p>Target folder permissions can now be customized to restrict user access.</p>	189583
Snowflake Target	Issue	<p>Writing CSV files to Snowflake target table would fail because the ODBC driver client version 2.18.3 was not installed.</p>	189602
Oracle Source	Issue	<p>Tasks would sometimes fail on Linux when replicating from an Oracle standby.</p>	189711

Component/Process	Type	Description	Ref #
PostgreSQL Source	Enhancement	Added version support up to PostgreSQL 10.6.	189753
		<div style="border: 1px solid gray; padding: 5px;"> <p>Note Do not use the newly supported versions in a production environment without first consulting your Attunity Account Manager.</p> </div>	
PostgreSQL Source	Issue	Tasks would fail with an "SQL_ASCII not found" error. Support for SQL_ASCII encoding was added.	189838
PostgreSQL Source	Issue	The following false positive message would be encountered when replicating from a PostgreSQL source: DDL Artifacts configured schema 'public' is different than the trigger's schema 'name'. DDLs will not be captured	189840
SAP Application (DB) with IBM DB2 for LUW backend, replicating to Log Stream.	Issue	When resuming a task that was stopped to add tables, no changes were applied to the target and the number of queued tables displayed in the monitor was negative.	190179

Component/Process	Type	Description	Ref #
Oracle Source -Attunity Log Reader	Issue	The task stopped applying changes after upgrading to Replicate 6.3 in CDB/PDB environment .	190365
ODBC Target	Issue	Applying updated CLOB columns to a SAP HANA target would sometimes fail with the following error (excerpt): datatype: lob type comparison: line 1 col 88674 (at pos 88673) [1022502] (ar_odbc_stmt.c:3809)	190683
Oracle Source - Attunity Log Reader	Issue	When the connection to ASM was interrupted during the parallel ASM read, Replicate would not automatically reestablish the connection.	191110
Metadata Manager	Issue	When a source column was changed to make it nullable, Replicate would try and alter the user-created target column as well. This would result in an infinite target reconnection loop as well as the following error: 00000380: 2019-04-15T14:37:30 [TARGET_APPLY]T: ORA-01440: column to be modified must be empty to decrease precision or scale [1022307] (oracle_endpoint_apply.c:1306)	191349

Component/Process	Type	Description	Ref #
IBM DB2 for LUW	Issue	In rare scenarios, the commit timestamp of a change would not be updated.	191391
Replicate Server	Issue	Due to a memory leak related to browser-initiated connection requests, the Replicate Server component would consume unusually high memory on Linux.	191465
SAP HANA Target	Issue	When applying changes to TIMESTAMP and DATE columns to SAP HANA target, the following error would be encountered: inconsistent datatype: NCLOB is invalid for function to_timestamp	191520
File Channel	Issue	When a new table was added to the local File Channel task while it was stopped, changes would not be captured until the task was resumed, stopped, and then resumed again.	191900
Notifications	Issue	When an "Any errors" notification was defined, an empty notification would wrongly be sent on Apply Conflicts (which falls under the "Any warning" category).	191969
Log Stream	Issue	After recovering from an environmental error, some changes would not be captured.	192000

Component/Process	Type	Description	Ref #
Bidirectional Task	Issue	The task would fail during its initial run when the source schema defined in the "Loopback prevention settings" was different from the target schema.	192442
PostgreSQL Source	Issue	When the task resumed from a recoverable error, some changes would not be captured.	192460
Oracle Source - Attunity Log Reader	Issue	A new Oracle redo log entry introduced in Oracle 12.2 (for RAC environments), would result in task errors with the following message: Cannot parse a Log Record entry: it contains an unexpected number of XX elements	192508
IBM DB2 for LUW Source	Issue	Attempting to start a task migrating from IBM DB2 for LUW 11.1 would fail with the following message: message SQL2650N An invalid parameter "piEndLRI" was passed to the "db2ReadLog" read log API. Reason Code = "17"	192596

Component/Process	Type	Description	Ref #
Oracle Target	Issue	<p>Homogeneous replication of XMLTYPE type columns exceeding 4000 bytes would result in the following non-recoverable error:</p> <pre>"ORA-01461: can bind a LONG value only for insert into a LONG column [1022307]"</pre>	192652
Log Stream	Issue	<p>In rare situations, the Log Stream target would confirm receipt of a Change Operation before it was written to the Log Stream file.</p> <p>Stopping the task before the Change Operation was written to the Log Stream file would sometimes result in the change not being applied to the target when the task was resumed.</p>	192690
IBM DB2 for LUW Source	Enhancement	<p>Added support for the Start processing changes from:Source change position: Run option.</p>	192693
IBM DB2 for iSeries	Issue	<p>Capturing changes from a TIMESTAMP column would result in the last digit being converted to a "0" on the target.</p>	192730

Component/Process	Type	Description	Ref #
IBM Informix Source	Issue	During CDC, data from locale zh_hk.utf8 (i.e. Chinese letters) would be applied to the target incorrectly as the internal parameter for manually setting the code page was not working properly.	192759
IBM DB2 for z/OS Source	Issue	The task would encounter an error when processing a partitioned table with a negative DBID value. Due to the resulting error, the table would be suspended.	193086
Microsoft SQL Server Source	Issue	After dropping a table from the source target, no changes were replicated on any table.	193284
Kafka Target	Issue	The task would fail when the source Primary Key/Unique Index columns contained a null value.	193622
PostgreSQL Source (Homogeneous)	Issue	Running a Full Load only task that included MONEY data with a fraction, would load an incorrect fraction to the target.	193852
Log Stream	Issue	Reloading a CDC only replication task would fail with an error.	193934

Component/Process	Type	Description	Ref #
Batch optimized apply - IBM DB2 for LUW target (Using ODBC)	Issue	Batch optimized apply would fail on IBM DB2 for LUW when one of the target columns was a CLOB.	194120
Microsoft SQL Server Source	Issue	Test connection would fail against Microsoft SQL Server 2005, because the existence of the is_cdc_enabled column in the system catalog tables was required. However, this column was only introduced in later versions.	194170
SAP Application (DB)source with Oracle backend.	Issue	During Change Processing, spaces at the end of a value in a STRING column would not be trimmed.	194433
Snowflake Target	Enhancement	Added the executeTimeout internal parameter for overriding the default statement timeout value.	194452
Log Stream	Issue	Partition information would not be retrieved when loading partitions from a Log Stream source.	194489
IBM Informix Source	Issue	When stopping a task in the middle of a transaction and then resuming the task, sometimes records would be missing from the target (as the task resumed from the next transaction instead of the uncommitted transaction).	194498

Component/Process	Type	Description	Ref #
Amazon Redshift Target	Issue	The task would sometimes fail to load the table due to creating duplicate records on the target.	194607
SAP Application (DB) Source	Issue	A maximum of 99 tables would be replicated to the target, even though more than 99 tables were selected for replication.	194766
IBM DB2 for LUW Source	Issue	When using Parallel Load, some of the partitions were not loaded to the target.	194940
SAP Application (DB) with Oracle backend	Issue	When processing UPDATES to cluster tables and setting a timeout longer than 5 minutes, Replicate would crash if the backend table UPDATES did not affect any SAP tables.	195146
Amazon Kinesis Target	Issue	In some cases, not all the data would be replicated to the target.	195225
Snowflake on AWS	Issue	INSERT and UPDATE operations would sometimes not be applied to the target.	195393
UI - Endpoints		When selecting Hortonworks Data Platform (HDP) from the endpoints list, Google Cloud for PostgreSQL would be displayed instead.	195529

Component/Process	Type	Description	Ref #
Log Stream	Issue	When the CDPOS table in SAP Application (DB) source is setup with a source endpoint that uses Log Stream, the column added to Snowflake target is NULL (instead of containing the correct value).	195607
Oracle Source - General	Issue	When the database included tables with many partitions and sub-partitions, starting the task would take an unusually long time.	195702
Informix Source - Sorter	Issue	In some cases, data loss occurs due to missing records in the log, resulting from changes in the stream position format.	195818
SAP Application Source with SAP HANA backend	Issue	When replicating long text, Chinese characters are displayed as '?' on the target.	195995
SAP HANA source, SAP source	Issue	The TDFORMAT prefix would be added to the STXL's CLUSTD field.	195995

Component/Process	Type	Description	Ref #
Oracle Source - Attunity Log Reader	Issue	<p>When replicating from Oracle 18.x Kafka, the task would fail with the following error:</p> <pre>Failed to produce kafka message</pre> <p>The issue was caused by a missing V_\$CONTAINERS permission on the Oracle database. The required permission was added to the Help.</p>	196042
Oracle Source - General	Issue	Tables with a large number of partitions would take a very long time to process.	196054
Scheduler	Issue	In rare cases, scheduled tasks would start afresh (i.e. with Full Load) instead of resuming.	196136
SAP HANA Source - Trigger-based CDC	Issue	<p>Due to an issue with the SELECT statement, the following error would sometimes occur during change capture (excerpt):</p> <pre>General error;339 invalid number: [6930] attribute value is not a number;exception 70006930: attribute value is not a number;</pre>	196689

Component/Process	Type	Description	Ref #
Microsoft SQL Server Source	Issue	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
Documentation	Issue	A redundant space in the command for configuring the IAM Roles for EC2 Access Type caused an error when the command was copied.	197428
Oracle Source - Attunity Log Reader	Issue	The AR_H_USER header on the target would be empty when the following mode was set in the Oracle database: <code>audit_sys_operations=false</code>	197255
Oracle Source - Attunity Log Reader	Issue	Capturing the Insert From Select operation into an Oracle Index Organized Table would sometimes result in missing operations.	197756
Oracle Source - Attunity Log Reader	Issue	Task would fail when the v\$containers permission was not granted, even though it wasn't needed for the selected source tables. Now, a warning will be issued, but the task will not fail.	197413