
Service Pack Release Notes for Version 6.4 SP19 (build 6.4.0.550)

These release notes provide details on the resolved issues and enhancements included in SP19. All service packs are cumulative, meaning that they include the fixes and enhancements provided in previous packs.

SP19 (build 6.4.0.550)

This section describes the resolved issues and/or enhancements in this service pack.

Resolved Issues

Component/Process	Description	Internal Issue #	Customer Case #
Oracle Source	<p>When the Expose NUMBER as endpoint option was set to FLOAT, the task would fail with the following error:</p> <pre>ORA-03115: unsupported network datatype or representation [1022307] (oracle_endpoint_load.c:2111)</pre>	RPT-24898	1914403

SP18 (build 6.4.0.536)

This section describes the resolved issues and/or enhancements in this service pack.

Resolved Issues

Component/Process	Description	Internal Issue #	Customer Case #
Amazon Kinesis Target	<p>In a Kinesis environment with 200 streams or more, a "Rate Limit Exceeded" error would be encountered when clicking the "Browse" button to select a specific stream.</p>	RPT-24557	1931327
Google Dataproc Target	<p>After upgrade, Full Load would sometimes fail with the following error:</p> <pre>Failed to find descriptor for '114'</pre>	RPT-24138	1910952

SP17 (build 6.4.0.534)

This section describes the resolved issues and/or enhancements in this service pack.

Resolved Issues

Component/Process	Description	Internal Issue #	Customer Case #
Amazon Redshift Target	<p>Full Load replication in Parallel Load mode would sometimes fail with the following error:</p> <pre>The specified S3 prefix XXXXX does not exist</pre>	RPT-24355	1912124
Hadoop target	<p>When periodically running a Full Load task with the Archive and Create option and with the Target storage format set to ORC, the data in the target tables would remain the same, even though the source data had changed.</p>	RPT-23448	1886453

SP16 (build 6.4.0.532)

This section describes the resolved issues and/or enhancements in this service pack.

Enhancements

Component/Process	Description	Internal Issue #	Customer Case #
Kafka Target	<p>UPDATE and DELETE operations performed on source tables with non-nullable LOB columns would fail on Kafka target with the following error (excerpt):</p> <pre>Got unexpected null value in column.</pre>	RPT-23965	1905686
Amazon Kinesis Data Streams Target	<p>When using the endpoint [Browse] button to browse for a specific stream, if more than 100 streams were present on Amazon Kinesis Data Streams, no streams would be listed in the pop-up window and the following error would be returned (excerpt):</p> <pre>Failed to list the available streams on Amazon Kinesis Data Streams</pre> <p>When the stream name was entered manually (instead of using the Browse button), the task would fail.</p>	RPT-23970	1907700
Replicate General	<p>The original list of signal types that Replicate was able to capture in order to generate a core dump has been expanded to support many more signals such as SIGALRM, SIGBUS, SIGFPE, SIGHUP, SIGILL, SIGINT, SIGIO, and so on.</p>	RPT-22966	1833777

SP15 (build 6.4.0.529)

This section describes the resolved issues and/or enhancements in this service pack.

Enhancements

Component/Process	Description	Internal Issue #	Customer Case #
Hadoop target endpoint	Added a new internal parameter that allows the ODBC statement timeout length to be extended during periods of heavy load.	RPT-23792	1900365

SP14 (build 6.4.0.528)

This section describes the resolved issues and/or enhancements in this service pack.

Resolved Issues

Component/Process	Description	Internal Issue #	Customer Case #
Replicate General	Added an option to catch SIGSEGV and SIGABRT signal types and generate a core dump.	RPT-22966	1833777
Google Cloud Storage Target	Tables would sometimes not be loaded to the target during Full Load. To resolve this issue, the recovery mechanism was improved to handle connection errors more efficiently, and logging was enhanced to provide more details when set to "Trace".	RPT-23201	1872237

Component/Process	Description	Internal Issue #	Customer Case #
Microsoft SQL Server Source	During Change Processing, if the endpoint was set with Replicate has file-level access to the backup log files , the task would fail when parsing a rare log operation.	RPT-23788	1900340
Amazon Kinesis Target	Due to the time spent checking for new messages, batches would take several seconds to close, resulting in performance issues. The issue was resolved by streamlining the procedure used to check for new message.	RPT-23757	1871866

SP13 (build 6.4.0.517)

This section describes the resolved issues and enhancements in this service pack.

Resolved Issues

Component/Process	Description	Internal Issue #	Customer Case #
SAP Application Source with an IBM DB2 for iSeries or a Microsoft SQL Server backend	The task would stop unexpectedly when capturing DD02L events in the CAPTURE TRACE mode.	RPT-23013	1817375
Oracle Source - Attunity Log Reader	When using Replicate on Windows to read NFS-mounted redo logs, the task would encounter caching issues resulting in CDC slowing down/stopping. The issue was resolved using an internal property.	RPT-23141	1871875

SP12 (build 6.4.0.515)

This section describes the resolved issues and enhancements in this service pack.

Resolved Issues

Component/Process	Description	Internal Issue #	Customer Case #
Kafka target	A NULL value being replicated to a non-nullable column would result in task failure instead of a data error.	RPT-22996	1867042
File Target	When code page conversion failed the task would enter recovery resulting in duplicate rows being written to the file target. Now, when code conversion fails, an appropriate warning will be issued.	RPT-22937	1851234
Amazon Kinesis Target	Incorrect record sequence numbers and transaction counters on the target would result in failure of the consuming application.	RPT-23331	1878784
Replicate Server	High memory usage was encountered when repeatedly connecting to Replicate Server on Linux.	RPT-23237	1755802
IBM DB2 for iSeries Source	If a schema was dropped while the task was stopped, the task would fail instead of continuing with a warning.	RPT-23012	1866217
Replicate Sorter	In very rare cases (e.g. if the task started and then stopped immediately), changes would sometimes not be captured.	RPT-23327	1871235

Component/Process	Description	Internal Issue #	Customer Case #
SAP Application (DB) with Microsoft SQL Server backend	During Full Load, a "Tables not found" error would be encountered when trying to replicate the dictionary tables DD02L/DD03L, even though the tables were present.	RPT-23006	196469
Oracle Source	CDC task would fail due to new permissions requirements that were not needed for the task. An option to disable the required permissions (Grant SELECT ON V_\$DATABASE_INCARNATION and Grant SELECT ON GV_\$TRANSACTION) via an internal parameter was provided.	RPT-22821	196699
SAP HANA Source - Trigger-based CDC	The INDX column in the Replicate <code>cdc_changes</code> table was defined as LONG instead of BIGINT. This would cause Replicate to stop capturing changes when it reached the limit allowed by the LONG data type. To resolve this issue, the INDX column data type has been changed to BIGINT.	RPT-23101	192159

SP11 (build 6.4.0.502)

This section describes the resolved issues and enhancements in this service pack.

Resolved Issues

Component/Process	Description	Internal Issue #	Customer Case #
DB2 z/OS Source	Capture process could not resume the task after a table was dropped in the previous run.	RPT-22552	163867

Component/Process	Description	Internal Issue #	Customer Case #
Oracle Source - General	When performing a lookup from a LOB column, a memory leak would occur when a Unique Index column contained at least one NULL value.	RPT-22527	196758
DB2 LUW SOURCE, SAP (DB)	Out-of-row LOB values would not be recognized when the LONG IN clause was defined in the TABLE CREATE statement.	RPT-21884	193522
Microsoft SQL Server Target	The log message "Failed to load SQL Server Native Client 11.0" was printed to the log at Trace severity level rather than at Error level.	RPT-22744	197056
Microsoft SQL Server Source	When two tasks with a Microsoft SQL Server source endpoint configured to use the direct-read method were run in parallel, one of the tasks would fail.	RPT-22618	197107
Documentation	A redundant space in the command for configuring the IAM Roles for EC2 Access Type caused an error when the command was copied.	RPT-22714	197428
SAP HANA Source - Trigger-based CDC	<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>	RPT-22698	196689

Component/Process	Description	Internal Issue #	Customer Case #
Change Tables	When writing to Change Tables in Transactional apply mode, the TRANSACTION_ID column would contain non-printable characters (binary zeros instead of character zeros).	RPT-22402	196642
SAP Application (DB) with an Oracle backend using Attunity Log Reader	In rare cases, a memory issue would occur when capturing changes from BLOB columns when the task was set to limited LOB support.	RPT-22060	195803
Microsoft Azure Data Warehouse Target	Sometimes, when an error was encountered on the target, the external table would not be cleaned up. This would cause the task to fail on recovery as it could not create an external table (due to the previously created table already existing).	RPT-22839	1777641
IBM DB2 for LUW Source	Due to an issue with the DB2LUW Deferred Constructor, records with out-of-row LOBs would sometimes be ignored.	RPT-22954	1853736
IBM DB2 for iSeries Source	In rare cases, the metadata connection would time out during Full Load (after a large table was loaded). This would result in the suspension of all tables that had not been loaded to the target.	RPT-22941	1830470
Replicate server	Due to memory not being freed when a connection was closed, Replicate server memory consumption would be unusually high.	RPT-22939	1853355

Enhancements

Component/Process	Description	Internal Issue #	Customer Case #
SAP Database	Improved handling of large LOB data types.	RPT-22448	193522
Oracle Source	The Oracle client version will now be printed to the Replicate logs.	RPT-22796	N/A

SP10 (build 6.4.0.485)

This section describes the resolved issues and enhancements in this service pack.

Resolved Issues

Component/Process	Description	Internal Issue #	Customer Case #
Oracle Source - General	Tables with a large number of partitions would take a very long time to process.	RPT-22118	196054
Informix Source, Sorter	In some cases, data loss would occur in production due to missing records in the log caused by changes in the stream position format.	RPT-22009	195818
Oracle Source - Attunity Log Reader	In rare cases, while running Full Load, the task would fail with the following error: <code>Incorrect numeric value in the column XYZ is changed to 0</code>	RPT-21970	195696
Oracle Source - General	When the database included tables with many partitions and sub-partitions, starting the task would take an unusually long time.	RPT-22145	195702

Component/Process	Description	Internal Issue #	Customer Case #
SAP Application (DB) source with Oracle backend	The BLOB data type would be missing from the data record.	RPT-21972	195767
		RPT-22059	195911
Log Stream	When running a task while using Log Stream between the source and target, the table IDs metadata for SAP tables would not specify the same IDs in the Replication task.	RPT-21922	195607
SAP Application Source with SAP HANA backend	When replicating long text, Chinese characters would be displayed as '?' on the target.	RPT-22104	195995
SAP HANA Log-Based Source	<p>After running CDC scripts and comparing the counts between source and target, some rows would be missing on the target.</p> <p>In addition, when running the CDC script multiple times without stopping the task, a Task stopped abnormally error would be displayed but recover after a few seconds and start applying the changes.</p>	RPT-20986	192412
SAP Application (DB) source with Oracle backend	<p>In rare cases, when replicating OLTP compressed rows, the following error would be encountered:</p> <pre>E: The field 'fieldname' doesn't exist in the CDC record for table 'Business Group'. 'tablename'</pre>	RPT-21804	195171
Scheduler	In rare cases, scheduled tasks would start afresh (i.e. with Full Load) instead of resuming.	RPT-22180	196136
REST API	The Attunity Replicate Server service would keep stopping.	RPT-22271	196187

Component/Process	Description	Internal Issue #	Customer Case #
SQLite	Fixed a critical SQLite vulnerability.	RPT-22372	N/A
Microsoft SQL Server to Log Stream Target	The task would crash when capturing changes from Microsoft SQL Server using the endpoint's Replicate has file-level access to the backup log files option.	RPT-22305	196645
Oracle Source	An "interval is invalid" error would be encountered when trying to insert an INTERVAL data type with precision "0" to the target.	RPT-22368	196643
Oracle Source - Attuniy Log Reader	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>	RPT-22557	197255

Enhancements

Component/Process	Description	Internal Issue #	Customer Case #
SAP HANA source, SAP source	Added an option using an internal parameter to remove TDFORMAT prefix from the STXL's CLUSTD field.	RPT-22179	195995

SP09 (build 6.4.0.457)

This section describes the resolved issues and enhancements in this service pack.

Resolved Issues

Component/Process	Description	Internal Issue #	Customer Case #
IBM DB2 for LUW Source	Out-of-row LOB values would not be processed correctly.	RPT-21884	193522
SAP Application (DB)	<p>Extra logging was added and a dump record will be created when the following error is encountered:</p> <pre>E: The field 'fieldname' doesn't exist in the CDC record for table 'Business Group'. 'tablename'</pre>	RPT-21804	195171
Snowflake on AWS	INSERT and UPDATE operations would sometimes not be applied to the target.	RPT-21969	195393
UI - Endpoints	When selecting Hortonworks Data Platform (HDP) from the endpoints list, Google Cloud for PostgreSQL would be displayed instead.	RPT-21870	195529
Log Stream	During Full Load replication from SAP Application (DB) source, some of the rows in the in the CDPOS table's Extract Time and Date columns would not be replicated to the target.	RPT-21922	195607

SP08 (build 6.4.0.446)

This section describes the resolved issues and enhancements in this service pack.

Resolved Issues

Component/Process	Description	Internal Issue #	Customer Case #
SAP Application (DB)source with Oracle backend.	During Change Processing, spaces at the end of a value in a STRING column would not be trimmed.	RPT-21654	194433
Console UI	The task Settings button would be displayed without a label.	RPT-20907	192848
SAP HANA source (Log-based CDC)	During Change Processing, new source records would sometimes be missing from the target.	RPT-20986	192412
Oracle Source	Replicate would wrongly identify ORA-14402 as an environment error instead of a data error.	RPT-21716	193537
IBM Informix Source	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).	RPT-21634	194498
Oracle - Attunity Log Reader	In rare cases, when replicating OLTP compressed rows, Replicate performed DELETE operations instead of UPDATE.	RPT-21706	193700
SAP Application (DB) with Oracle backend	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.	RPT-21764	195146

Component/Process	Description	Internal Issue #	Customer Case #
Google SQL for MySQL Target	A syntax error would be encountered when inserting from the attrep_changes table if the target table contained a VARCHAR column.	RPT-21699	194699
IBM DB2 for LUW Source	When using Parallel Load, some of the partitions were not loaded to the target.	RPT-21721	194940
Amazon Redshift Target	In certain cases, some of the CSV files would be loaded to the wrong tables.	RPT-21777	194607
SAP Application (DB) Source	A maximum of 99 tables would be replicated to the target, even though more than 99 tables were selected for replication.	RPT-21638	194766
Microsoft SQL Server Source	If a source table was dropped and the task was then stopped and resumed, Replicate would stop capturing changes on all tables.	RPT-21812	193284

SP07 (build 6.4.0.425)

This section describes the resolved issues and enhancements in this service pack.

Resolved Issues

Component/Process	Description	Internal Issue #	Customer Case #
Microsoft SQL Server Source	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.	RPT-21480	194170

Component/Process	Description	Internal Issue #	Customer Case #
Amazon Redshift Target	The task would sometimes fail to load the table due to creating duplicate records on the target.	RPT-21620	194607
Kafka Target	The task would fail when the source Primary Key/Unique Index columns contained a null value.	RPT-21282	193622

Enhancements

Component/Process	Description	Internal Issue #	Customer Case #
Snowflake Target	Added the executeTimeout internal parameter for overriding the default statement timeout value.	RPT-21533	194452

SP06 (build 6.4.0.404)

This section describes the resolved issues and enhancements in this service pack.

Resolved Issues

Component/Process	Description	Internal Issue #	Customer Case #
Log Stream	When an incorrect storage path was defined in the Log Stream target endpoint, the following message would be incorrectly included in the logging: <code>ORA-01031: insufficient privileges</code>	RPT-21435	192241
Batch optimized apply - IBM DB2 for LUW target (Using ODBC)	Batch optimized apply would fail on IBM DB2 for LUW when one of the target columns was a CLOB.	RPT-21457	194120

Enhancements

Component/Process	Description	Internal Issue #	Customer Case #
Microsoft Azure ADLS	Added support for Azure Data Lake Storage (ADLS) Gen2 storage.	RPT-21082	N/A
Microsoft Azure Databricks	Added support for Azure Data Lake Storage (ADLS) Gen2 storage.	RPT-21084	N/A
Endpoints	The Microsoft Azure Data Warehouse endpoint was renamed to Microsoft Azure SQL Data Warehouse .	RPT-21366	N/A

SP05 (build 6.4.0.396)

This section describes the resolved issues and enhancements in this service pack.

Resolved Issues

Component/Process	Description	Internal Issue #	Customer Case #
IBM DB2 for z/OS Source	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.	RPT-21054	193086
IBM DB2 for z/OS Source	Tables with BRF and RRF rows were handled incorrectly, resulting in the following error: <code>Can't match columns to record data</code>	RPT-21311	187565

Component/Process	Description	Internal Issue #	Customer Case #
IBM Informix Source	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.	RPT-20962	192759
PostgreSQL Target	Excessive memory would be consumed when using the "Load" command to load CSV files to the target.	RPT-20704	192243
Sap Application with IBM DB2 for LUW backend	When the backend was a non-unicode IBM DB2 for LUW database, the following error would be encountered when starting the task: RetCode: SQL_ERROR SqlState: 53090 NativeError: -873 Message: [IBM][CLI Driver][DB2/AIX64] SQL0873N Objects encoded with different encoding schemes cannot be referenced in the same SQL statement. SQLSTATE=53090	RPT-21063	193156
Metadata Manager	Bidirectional task would fail on initial run with the following error: attrep_loopback_prevention' internal table has an error	RPT-21353	193926
Log Stream	The \$AR_H_USER variable would return a null value in a Log Stream Staging task.	RPT-21034	193092
Log Stream	Reloading a CDC only replication task would fail with an error.	RPT-21367	193934

Enhancements

Component/Process	Description	Internal Issue #	Customer Case #
IBM DB2 for z/OS Source	Following the initial positioning enhancement in SP03 - RPT-20854 , support for backward compatibility was added to enable customers who are unable to upgrade R4Z immediately, to continue working with their current configuration.	RPT-21234	N/A
Sap Application Source	SAP TIME value 240000 would be incorrectly converted to 230000. The issue was resolved using an internal parameter.	RPT-21240	N/A
IBM DB2 for LUW Source	Added an internal parameter that converts a 24:xx:xx time_format string to a user-specified time format (hh:mm:ss).	RPT-21052	N/A
General (Validation)	Added support for performing transformations on column names with spaces, even if the column doesn't already exist on the target.	RPT-21332	N/A
MySQL Source	Added support for capturing XA transactions from MySQL 5.7.	RPT-21002	192726

Component/Process	Description	Internal Issue #	Customer Case #
Endpoint Failover	<p>Added an SDK method for overriding endpoint settings with those of another endpoint (to support failover scenarios).</p> <p>The method name differs according to the SDK type:</p> <ul style="list-style-type: none"> » REST SDK: <pre>AEMReconfigureEndpointNoWait</pre> » .NET SDK: <pre>ReconfigureEndpointNoWait</pre> » Python SDK: <pre>reconfigure_endpoint_no_wait</pre> <p>For usage instructions, please refer to the <i>AEM Developer's Guide</i>.</p>	RPT-21396	189608

Note The new method is currently supported with the Oracle source endpoint only.

SP04 (build 6.4.0.371)

This section describes the enhancements and resolved issues in this service pack.

Resolved Issues

Component/Process	Description	Internal Issue #	Customer Case #
SAP Application Source with SAP HANA backend	Changes from the Long Text STXL table would not be captured properly.	RPT-21029	193057
Hadoop Target	On rare occasions, an error would occur when using Sequence format with GZIP compression.	RPT-20829	188415
Oracle Source - Attunity Log Reader	<p>A new Oracle redo log entry introduced in Oracle 12.2 (for RAC environments), would result in task errors with the following message:</p> <pre>Cannot parse a Log Record entry: it contains an unexpected number of XX elements</pre>	RPT-21128	192508
Oracle Homogeneous	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.	RPT-21050	185286
Sap Application	<p>SAP TIME value 240000 would be incorrectly converted to 230000.</p> <p>The issue was resolved using an internal parameter.</p>	RPT-21192	193522

Enhancements

Component/Process	Description	Internal Issue #	Customer Case #
IBM DB2 for iSeries Source	Added support for retrieving system column names via an internal parameter.	RPT-19323	190691
Amazon Redshift Target	Added support for uploading multiple files (up to 100) concurrently for improved performance.	RPT-20916	N/A
IBM DB2 for iSeries Source	Added support for capturing changes via UDTF rather than calling display_journal.	RPT-21088	N/A
Endpoints	Added support for Google SQL Cloud for MySQL source.	RPT-20623	N/A

SP03 (build 6.4.0.353)

This section describes the enhancements and resolved issues in this service pack.

Resolved Issues

Component/Process	Description	Internal Issue #	Customer Case #
Microsoft SQL Server Source	During CDC, the DATETIMEOFFSET data type would be replicated with an incorrect precision value.	RPT-19716	191231
SAP Application (DB) Source with Microsoft SQL Server Backend	When replicating the BSEG cluster table, a Primary Key violation would be encountered at the start of Full Load.	RPT-20613	192270

Component/Process	Description	Internal Issue #	Customer Case #
SAP Application (DB) Source with Oracle Backend	<p>When trying to perform Full Load of a newly added table, the following error would be encountered (excerpt):</p> <pre>Failed sending row to table...String data, right truncation.</pre>	RPT-20648	190840
File Channel	<p>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.</p>	RPT-20709	191900
Oracle Source - Attunity Log reader	<p>If the connection to ASM failed <i>while</i> the redo log was being opened, the connection would not be reestablished.</p>	RPT-20713	192377
IBM DB2 for LUW Source	<p>In rare scenarios, the commit timestamp of a change would not be updated.</p>	RPT-20499	191391
Oracle Source - Attunity Log reader	<p>After upgrading Oracle CDB/PDB from 12.1 to 12.2, changes would no longer be captured.</p>	RPT-20747	192522
Oracle Source - Attunity Log reader	<p>When using redo log version 18 on Oracle 18c, the following parsing error would be encountered:</p> <pre>Cannot parse a Log Record, continuation. Problematic entry #0. Offset 0. Record data length 260</pre>	RPT-20786	190618

Component/Process	Description	Internal Issue #	Customer Case #
Log Stream	<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>	RPT-20796	192690
PostgreSQL Source	When the task resumed from a recoverable error, some changes would not be captured.	RPT-20701	192460
IBM DB2 for iSeries	Capturing changes from a TIMESTAMP column would result in the last digit being converted to a "0" on the target.	RPT-20891	192730
Oracle Source - Attunity Log Reader	<p>In rare situations, the following warning would be encountered when capturing changes from Oracle, even though the replication did not include encrypted columns:</p> <p>Redo log contains both encrypted and unencrypted events for table...only unencrypted events will be captured.</p>	RPT-19666	189120
ARC Endpoints	<p>Reloading an individual table via the console UI would fail with the following error:</p> <pre>[TASK_MANAGER]W: Table '(null)'. 'TABNAME' cannot be found. (replicationtask_cmd.c:838)</pre>	RPT-20501	191833

Component/Process	Description	Internal Issue #	Customer Case #
Sap Application (DB)	SAP TIME value 240000 would be incorrectly converted to 230000. The issue was resolved using an internal parameter.	RPT-20922	191986
IBM DB2 for z/OS	DDLs consisting of two-record updates would sometimes result in Table and Change Tables being dropped from the target.	RPT-20856	191174
Sap Application (DB) with IBM DB2 for LUW	The task would stop unexpectedly during Full Load when encountering a VARDATA column with an empty value in clustered tables.	RPT-20925	192537
Oracle Source	Due to an issue with the supplemental logging verification process, a Change Processing task defined with a manipulation to change an index in a table with a LOB would fail. The failure was due to supplemental logging not being defined for the source table index. Solution: Define the supplemental logging on the changed key segments manually, or reload the table.	RPT-20862	192727
IBM DB2 for LUW Source	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"	RPT-20824	192596

Component/Process	Description	Internal Issue #	Customer Case #
Microsoft SQL Server Target	Mapping of the TIME data type was changed from TIME (7) to TIME for all Microsoft SQL Server-based targets.	RPT-20539	N/A
Notifications	When an "Any errors" notification was defined, an empty notification would wrongly be sent on Apply Conflicts (which falls under the "Any warning" category).	RPT-20767	191969
Oracle Target	The task would fail with the following error (excerpt): <code>Failed checking if index name (attrep_changes86D165) is unique</code>	RPT-20965	186051
UI	Opening multiple task tabs and then trying to select one of the open tabs would always redirect to the first tab that was opened.	RPT-20983	192286 192834 192515 192459
Oracle Target	Homogeneous replication of XMLTYPE type columns exceeding 4000 bytes would result in the following non-recoverable error: <code>"ORA-01461: can bind a LONG value only for insert into a LONG column [1022307]"</code>	RPT-20953	192652
MySQL Source	When reassuming a task, if the stream_ position contained the name of an earlier binlog for the table_map, the task would wrongly resume from the old binlog instead of the current binlog, and eventually fail.	RPT-21004	192882

Enhancements

Component/Process	Description	Internal Issue #	Customer Case #
Microsoft Azure Event Hubs Target	Support for TLS 1.2 was added.	RPT-20684	192090
IBM DB2 for z/OS	Implemented syslgrnx as part of the initial positioning to avoid the potential "Start from timestamp" issue of overshooting the desired start time.	RPT-20854	N/A
Microsoft SQL Server Source	Added support for using a Replicate non-sysadmin user to access an AlwaysOn configuration.	RPT-20621	N/A
MemSQL Target	An option was added to create the target tables as columnstores.	RPT-20771	N/A
Kafka Target	Improved performance when using Avro format with tables containing large numbers of columns.	RPT-20959	N/A
Microsoft Azure Databricks	Added support for Microsoft Azure Databricks with Blob storage only.	RPT-21035	N/A
Google Cloud for PostgreSQL Source	Added support for Google Cloud for PostgreSQL Source (Full Load only).	RPT-20622	N/A
Teradata Target	<p>Added support for Teradata target version 16.20.</p> <p>Note that replicating to this version requires installing ODBC driver version 16.2 and setting the driver name as an internal parameter (Teradata Database ODBC driver 16.20).</p>	RPT-20915	N/A
Hadoop Target	Added support for Cloudera 6.1.	RPT-21089	N/A

SP02 (build 6.4.0.316)

This section describes the enhancements and resolved issues in this service pack.

Resolved Issues

Component/Process	Description	Internal Issue #	Customer Case #
Metadata Manager	<p>When a source column was changed to make it nullable, Replicate would try and alter the user-created target column as well.</p> <p>This would result in an infinite target reconnection loop as well as the following error:</p> <pre>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)</pre>	RPT-19596	191349
SAP HANA - Log-Based Source	Capturing a large number of changes would sometimes result in excessive memory consumption.	RPT-20556	
Log Stream	After recovering from an environmental error, some changes would not be captured.	RPT-20581	192000
SAP Transports	<p>A missing table (DD16S) in newer versions of SAP HANA would cause the import of the first transport file to fail with the following error:</p> <pre>Type 'DD16S' is unknown.</pre>	RPT-20566	191694

Component/Process	Description	Internal Issue #	Customer Case #
SAP Transports	When the Gold Client solution and the Replicate for SAP solution were both installed on the same system and using same client, any modifications made by via the ZGOLD interface or the ZR4SAP interface would also affect the other solution (as they shared the same configuration tables).	RPT-20569	189187
Hadoop Target	Reloading a task during Full Load would fail due to memory leak when using gzip compression.	RPT-20573	190396

SP01 (build 6.4.0.311)

This section describes the enhancements and resolved issues in this service pack.

Resolved Issues

Component/Process	Description	Internal Issue #	Customer Case #
Log Stream with IBM DB2 for iSeries Source	<p>Resumed tasks would fail with the following error:</p> <pre>Data Record has unmatching count for ids(1), values(1), flags(0), bookmark ids(0) or bookmark values(0) [1020101] (ar_cdc_channel.c:985)</pre> <p>This issue was resolved and extra logging was added.</p>	RPT-19481	189498
Oracle Source - Attunity Log Reader	When the connection to ASM was interrupted during the parallel ASM read, Replicate would not automatically reestablish the connection.	RPT-19472	191110

Component/Process	Description	Internal Issue #	Customer Case #
SAP Application Source with Oracle backend to Microsoft SQL Server Target	Source date columns are being inserted to the target as NULL values. Extra logging was added to help troubleshoot the issue.	RPT-19426	190842
Log Stream	Errors would occur intermittently with various replication tasks reading from log stream staging tasks. Each time the errors occurred the replication task would stop and then resume automatically, causing the Processed Changes counter to be reset.	RPT-19398	189798
Log Stream	Reloading a stopped Log Stream Staging task would sometimes cause the Replication task to consume 100% CPU.	RPT-19296	190646
IBM DB2 for z/OS	Attempting to resume a task from timestamp would sometimes fail with the following error: <code>Position by time exceeded duration limit</code> The positioning timeout is now dependent on the endpoint's execution timeout, which can be modified using R4Z 6.4.0.281, which is included in this Service Pack.	RPT-19647	191243
File Channel + SAP Sybase ASE Source	The remote File Channel task would stop unexpectedly when an ALTER COLUMN operation was performed on the source table.	RPT-19592	190617

Component/Process	Description	Internal Issue #	Customer Case #
ODBC Target	Applying updated CLOB columns to a SAP HANA target would sometimes fail with the following error (excerpt): <pre>datatype: lob type comparison: line 1 col 88674 (at pos 88673) [1022502] (ar_odbc_stmt.c:3809)</pre>	RPT-19632	190683
Google Cloud Storage Target	Newly added folders would not be visible when browsing the bucket via the endpoint "Browse" button.	RPT-19518	188660
Google Dataproc Target	Newly added folders would not be visible when browsing the bucket via the endpoint "Browse" button.	RPT-19519	189196
Oracle Source - Attunity Log Reader	In rare situations, UPDATES would not be captured as a result of validation being erroneously performed on "SPLIT UNDO" operations.	RPT-19664	N/A
Microsoft SQL Server	When trying to resolve a partition ID that no longer exists, CDC would fail with the following error: <pre>Failed in resolving object ID by partition ID</pre>	RPT-19387	188384
Hadoop Target	Upload of the CSV file to HDFS would fail because the file size was zero.	RPT-19033	188762
Bulk Apply	UPDATES would fail in Batch Optimized Apply mode with MERGE syntax when the Unique Index contained nullable columns on the target.	RPT-19735	N/A
SAP HANA Target	Columns of BLOB and CLOB data types would be truncated to the default size (8K) when loaded to the target.	RPT-19750	N/A

Component/Process	Description	Internal Issue #	Customer Case #
Microsoft SQL Server to Kafka Target	<p>Out of row data would not be converted to UTF-8 format.</p> <p>This would result in the following error (excerpt):</p> <pre>Failed to convert data to json (value_to_plain_json), status = 12: Failed to set json value for field 'message'. Failed to set json value for field 'data'.</pre>	RPT-19400	186503
Microsoft SQL Server	Dues to incorrect handling of zero length columns, the data from one source column would be applied to multiple target columns during CDC.	RPT-19703	187559
Microsoft SQL Server to Kafka Target	The BEFORE-IMAGE for tables altered with ADD COLUMN would sometimes contain corrupted data.	RPT-19586	186503
Replicate Server	Due to a memory leak related to browser-initiated connection requests, the Replicate Server component would consume unusually high memory on Linux.	RPT-19759	191465

Enhancements

Component/Process	Description	Internal Issue #	Customer Case #
SAP HANA Source - Log-based CDC	Added support for capturing "Insert from select" operations and improved logging.	RPT-19377	N/A

Component/Process	Description	Internal Issue #	Customer Case #
Microsoft SQL Server Source	<p>Unloading data from the source exceeded the default timeout of 15 minutes.</p> <p>An internal parameter was added that enables users to set the desired timeout.</p>	RPT-19548	191149
Parallel Load	<p>Parallel Load support was added for the following targets:</p> <p>Snowflake on Azure, Snowflake on AWS, and Microsoft Azure Data Warehouse.</p>	RPT-19752	N/A
Parallel Load	Parallel Load support was added for the Amazon RDS for Microsoft SQL Server source.	RPT-20536	N/A
Change Processing	Added support for using the Apply batched changes to multiple tables concurrently option with the Amazon Redshift, and Microsoft Azure SQL Database target endpoints.	RPT-19751	N/A
Pivotal Greenplum Target	Added support for Pivotal Greenplum versions 5.7 and 5.16.	RPT-19782	N/A
Snowflake on Azure/AWS Target	Improved consistency when loading several tables to the target concurrently.	RPT-19763 RPT-19506	191623 191038
Logging	INSERT information was added to the TRACE logging when using the Apply batched changes to multiple tables concurrently option.	RPT-19760	N/A
Sybase ASE	Added support for Sybase ASE source on Linux.	RPT-19636	N/A

Component/Process	Description	Internal Issue #	Customer Case #
PostgreSQL Target Versions	Added support for: <ul style="list-style-type: none"> » PostgreSQL versions 10.7,11.1, 11.2 » Amazon RDS for PostgreSQL version 11.1 » Google Cloud for PostgreSQL version 11.1 	RPT-20533	N/A
PostgreSQL Source Versions	Added support for: <ul style="list-style-type: none"> » PostgreSQL versions 10.7,11.1, 11.2 » Amazon RDS for PostgreSQL version 11.1 	RPT-20534	N/A