Attunity Replicate 6.3 Release Notes - December 2018

Replicate 6.3 introduces several new features and important enhancements including Log Stream, upload optimization for Hadoop and Amazon S3, custom message format for Kafka, and improved AlwaysOn support for Microsoft SQL Server source.

> **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

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

Replicate Server Behind a Proxy Server or on Docker

Due to an enhancement to the XSRF token (which addresses the "Cross-Site Request Forgery Prevention Cheat Sheet" security breach), placing Replicate Server behind a proxy or creating it on Docker will prevent the Replicate and AEM web consoles from being able to establish a connection to Replicate Server (resulting in an HTTP 307 redirect error).

Note that this will only happen if the proxy server and Replicate server are configured with different ports.

The workaround is to disable the usage of XSRF as follows:

1. Edit the file **repctl.cfg** (locates in <PRODUCT_DIR>\bin) and add the following line:
   `"enable_xsrf_protection":false`
2. Save the file.
3. Restart the **Attunity Replicate Server** service.
   A warning will be printed to the log file that xsrf protection is disabled.

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 or 6.0 configured with the "data" folder in a non-default location, you must add the following parameter to the upgrade command:

`data=existing_replicate_data-directory`

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

`data=/opt/mydatadir/ rpm -U[vh] areplicate-6.2.0-102.x86_64.rpm`

Kafka Target Endpoint

JSON Message Structure Changes - Backward Compatibility Support

Replicate 5.2 introduced changes to the JSON message structure. During upgrade, if Replicate detects any tasks with a Kafka endpoint, it will enable the internal parameter `useOldJsonVersion`. This will allow the tasks to continue running using the old JSON

format. Customers who wish to start using the new JSON structure need to stop the task, clear the `useOldJsonVersion` internal parameter (in the **Advanced** tab), and then resume the task.

> **Note** Upgrade support for tasks using the old JSON structure will be discontinued from Replicate 6.4.

Avro Logical Types - Backward Compatibility Support

Replicate 6.2 introduced support for Avro Logical Data Types, with new mappings for some data types. On upgrade, the internal parameter `useOldUintMapping` will be automatically set in order to maintain backward compatibility with the previous data type mappings.

Replicate-to-Avro Data Type Mappings Prior to Replicate 6.2:

UI4 --> INT

UI8 --> LONG

Replicate-to-Avro Data Type Mappings from Replicate 6.2:

UI4 --> LONG

UI8 --> STRING

If you wish to use the new mappings, simply disable this parameter.

Oracle Source Endpoint

Customers upgrading from versions earlier than Replicate 6.0 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.

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 versions prior to 6.1 is not supported. Consequently, customers with Replicate versions prior to 6.1 and 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*.

Replicating System Names with the New IBM DB2 for iSeries Endpoint

The IBM DB2 iSeries source endpoint now replicates tables based on their SQL names with unlimited length. If your IBM DB2 for iSeries database does not use SQL names, it's likely that you'll want to keep the replication based on system names.

To do this, enable the **Use table and schema system names** option in the **Advanced** tab.

File Target Endpoint

» From Attunity Replicate 5.5, the Replicate data types are included in the metadata file (DFM). After upgrading, customers with existing tasks that wish to use the new data types should stop the task, clear the `dfmVersion=1.0` internal parameter, and then restart the task.

» From Attunity Replicate 6.0, when the **Add metadata header with data types** option is enabled in the endpoint settings' **General** tab, Replicate data types will be added to the header. To allow customers who enabled this option to continue working with their existing setup after upgrade, the `syntax=FileSyntaxV1` internal parameter was added.

After upgrading, customers with existing tasks that wish to switch to the Replicate data types should:

1. Stop the task.
2. Clear the `syntax=FileSyntaxV1` internal parameter in the **Advanced** tab.
3. Restart the task.

Upgrading the Attunity Replicate Console Only

The following issue applies when upgrading from versions earlier than Replicate 6.0 *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.

Attunity Enterprise Manager (AEM) Compatibility

Replicate 6.3 is compatible with AEM 6.3 only.

New and Enhanced Features

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

Log Stream

The new Log Stream feature significantly enhances the replication process by constantly streaming changes captured from a selected source endpoint to a staging folder on Replicate Server.

This offers multiple benefits:

» When multiple CDC tasks read changes from large databases that share a single transaction log (e.g. Oracle and IBM DB2 for z/OS), the log is polled for only one target (the Log Stream target) instead of being polled separately for each target. This improves performance while greatly reducing the load on the source database and network.

» Changes are continuously written to the log stream regardless of the availability of the target endpoint. This is especially beneficial in situations where the original source transaction logs are no longer online or cannot be accessed efficiently. Additionally, this approach significantly reduces latency when the target endpoint becomes available again as the logs are read from the staging folder instead of from the source database.

» As the changes are constantly streamed to the Replicate server machine, the Replicate SLA that requires an extended log retention policy to be defined on the source in case the target cannot be accessed does not need to be enforced anymore. This allows you to reallocate valuable resources on the source while ensuring that log files will be up-to-date (on the Replicate Server machine) even if the target become unavailable.

» When replicating from a single source database to multiple targets, each target can have its own subset of the log stream data (tables, rows, etc.) and be started, stopped, and scheduled independently of the others.

Big Data Changes and Enhancements

Kafka Changes and Enhancements

Custom Message Format

When a task is defined with a Kafka target endpoint, you can specify a custom message format that will override the default Replicate message format. This may be useful if the consumer application needs to process the message in a particular format.

The custom message format can be defined at task level - in the Task Settings **Message Format** tab - and/or at table level (in the Table Settings **Message Format** tab). When it is defined at both task and table level, the message format defined for the table will take precedence over the message format defined for the task.

You can also specify a custom message key format when the Kafka target endpoint is defined accordingly (e.g. the message format is set to Avro). For full details, refer to the *Attunity Replicate Setup and User Guide*.

Other Kafka Enhancements

» **Support for Message Keys in Avro Format** - A new **Encode message key in Avro format** check box has been added to the **General** tab of the Kafka target endpoint settings. This option is only available when the message **Format** is set to **Avro**, **Publish** is set to **Publish data schemas to Confluent Schema Registry**, and the **Message Key** is *not* set to **None**.

» **Relocation of Fields** - The **Include Before-image in UPDATE messages** and **Include external Schema ID header** fields have been moved to the new **Message Format** tab (see Custom Message Format above) in the **Task Settings** and **Table Settings** windows. Note that during upgrade, these fields will be enabled automatically if they were originally enabled in the endpoint settings.

Improvements to Streaming Target Endpoints

New Namespace Field

A **Namespace** field, which acts as a unique identifier, has been added to the message. When working with the Kafka endpoint, the Namespace can be included/excluded or its value edited using the **Custom Message Format** feature described above.

New Headers for Ensuring Transactional Consistency

To alert consumer applications when transactions are ready for processing (and thereby preserve transactional consistency), the following headers were added to the data message:

» **transactionEventCounter** - The sequence number of the current operation in the transaction. This can be used to determine the order of operations within a transaction.

» **transactionLastEvent** - "true" indicates that it is the final record in the transaction whereas "false" indicates that not all of the records have been processed.

When working with the Kafka endpoint, the headers can be included/excluded using the **Custom Message Format** feature described above.

Improved Recovery Mechanism for Preserving Transactional Consistency

When a task is resumed or recovered, the task will now start/recover from the beginning of the transaction in order to preserve transactional consistency.

Manually Editing the Default Message Structure (Applies to all streaming endpoints *except* Kafka)

Should you wish to edit/exclude the Namespace value or exclude the transactional consistency headers, you need to export the task, edit the desired values (in the exported

JSON file), and then re-import the task. If you need assistance with performing this procedure, please contact Attunity Support.

Upload Optimization to Hadoop and Amazon S3

A new **File Uploads** tab has been added to the Task Settings. Clicking the **Optimize File Uploads** button in this tab will significantly improve performance when replicating to Amazon S3 and Hadoop. This feature is especially suited to organizations that have a high volume of data throughput and that are not subject to bandwidth limitations.

As well as optimizing the upload by transferring the files in parallel (instead of one-by-one), the new feature almost completely eliminates upload latency.

Support for Source Metadata in the Data Files

The **Add Metadata Header** option (which was previously available only for Hadoop) has now been added to the **Advanced** tab of the Amazon EMR and Microsoft Azure HDInsight target endpoint settings. When the target storage format is set to **Text**, you can optionally add a header row to the data files. The header row can contain the source column names and/or the intermediate (i.e. Replicate) data types.

Error Message Consolidation

Audit messages resulting from the same error have now been merged into a single message, thereby making it easier to identify and resolve task or server issues.

Specific Event IDs in Windows Event Log

In previous versions, all Replicate events/errors had a single Event ID (261). From Replicate 6.3, each event type now has its own Event ID.

For a full list of the Event IDs, refer to the section **Replicate Event IDs in Windows Event Log** in the *Attunity Replicate Setup and User Guide*.

Improvements to SAP Sybase ASE Source Endpoint

The following options have been added to the **Advanced** tab:

- » **Advance truncation point every (seconds):** The frequency with which to advance the secondary truncation point when reading changes from the SAP Sybase ASE log.
- » **Dump the log every (seconds):** How often to dump the parts of the log which Replicate has already read. The logs will be dumped according to the configuration of the SAP Sybase ASE database (i.e. archived, deleted, etc.).

Improvements to Microsoft SQL Server Source Endpoint

Enhanced AlwaysOn Support

» The new implementation eliminates the need to specify a backup replica

» Now supported on both Linux and Windows

» The enhancement maintains backward compatibility with the previous implementation

Compressed Backup Support

Replicate 6.3 now supports reading compressed backup logs directly from Microsoft SQL Server source. Note that this feature is currently supported on Replicate for Windows only.

Support for Applying Updates According to a Specific Condition

Starting from Replicate 6.3, you can define an expression that instructs Replicate only to apply UPDATEs when a user-defined condition has been met, for example, only when specific columns have changed. This is useful in situations when there are many updates in the source that the user has deemed not relevant for the target. With this feature, such 'irrelevant' updates are ignored.

For a usage example, please refer to section "Creating a Record Selection Condition for One or More Columns" in the *Attunity Replicate Setup and User Guide*.

Accessing the Previous Value of a Column in a Transformation

Starting from Replicate 6.3, you can define an expression that lets you access the previous value of an updated column in a transformation. This is only available when the before-image of a record is captured.

For a usage instructions, please refer to section "Creating a Record Selection Condition for One or More Columns" in the *Attunity Replicate Setup and User Guide*.

Revamped SAP Application Source Endpoint

Replicate 6.3 introduces a revamped SAP Application endpoint - **SAP Application (DB)** - which works exclusively with the database underlying the SAP Application. The revamped endpoint improves on the existing SAP Application endpoint in several major areas:

The revamped SAP Application source endpoint does not connect to the SAP Application and does not use RFC calls to resolve cluster and pool tables. This greatly reduces the footprint of the replication process on the SAP Application.

Additionally, the revamped SAP Application source endpoint intelligently analyzes changes to cluster and pool tables to find the minimal change between the old and new version of a row. Previously, an update to a row of a cluster or pool table was implemented by deleting all target rows that were associated with an updated source row, followed by inserting all

target rows now associated with the updated source row. Now, any change to a cluster or pool table row results in the minimal INSERT, UPDATE and DELETE operations on the target rows associated with updated source row. This greatly reduces the volume of operations on the target database while providing users with accurate information about the original operation that occurred on the source.

> **Note** In Replicate 6.3, the **SAP Application (DB)** endpoint is available only for SAP Applications working with an Oracle database. In future versions, additional platforms will be supported.

Snowflake Target with Microsoft Azure Blob Storage

From Replicate 6.3, customers replicating to Snowflake target can choose whether to store their files in Amazon S3 or in the newly supported Microsoft Azure Blob storage.

Other Enhancements

» **Oracle Source** - A **Number of concurrent ASM read threads** option has been added to the **Advanced** tab when working with Attunity Log Reader.

» **SSL connectivity has been added to the following endpoints:** MySQL source, Amazon RDS for MySQL source, PostgreSQL source, Amazon RDS for PostgreSQL source, PostgreSQL target, and Microsoft Azure Database for PostgreSQL target.

» **Microsoft Azure ADLS target** - Now supports Change Data Partitioning. For details about this feature, refer to the *Attunity Replicate User Guide and Setup*.

» **Parallel Load** - The Microsoft Azure SQL Database target endpoint is now supported.

New Endpoints and Version Support

This section lists new endpoint(s) as well as newly supported endpoint and platform versions.

New Platforms

Support has been added for Red Hat 7.4 and 7.5.

> **Note** If your Replicate Server is installed on Red Hat 7.4 and 7.5 and you wish to use Kerberos authentication with the Hadoop or Kafka endpoints, you need to rename the Kerberos libraries.
>
> The procedure for accomplishing this is described in the relevant chapters of the *Attunity Replicate Setup and User Guide*.

New Source Endpoints

Support has been added for the following source endpoints:
- Amazon RDS for PostgreSQL
- Amazon RDS for SQL Server
- Amazon RDS for MySQL

New Target Endpoints

Support has been added for the following target endpoints:
- Google Dataproc
- Google Cloud Storage

Newly Supported Endpoint Versions

- Oracle 18c Source
- Oracle 18c Target
- PostgreSQL Target 10.4, 10.5, and 10.6
- Kafka 1.1 and 2.0
- Cloudera 5.15 (Hadoop)

Deprecated Endpoint Versions

- PostgreSQL 9.0 (Target) - Support will be discontinued from Replicate 6.3.

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 6.1

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

For each Minor Release, only the latest Maintenance Release will be supported.

Based on the policy outlined above, Attunity announces the end of life of Replicate 6.1. Customers who need additional time to upgrade, can do so during the six-month period following the GA release of Replicate 6.3. Nevertheless, Replicate customers are encouraged to upgrade to Replicate 6.2 or above 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.

Resolved Issues and Customer Requested Enhancements

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

Component	Type	Description	Reference
Server	Issue	Replicate would crash when trying to retrieve monitoring data for tasks that were deleted from Replicate but still existed in AEM	N/A
IBM DB2 for z/OS Source	Issue	Resume from timestamp would not work when a large Result set size (1024 MB or greater) was specified.	N/A
Oracle Target	Issue	When int32 was a negative value, the task would fail with the following error: "Failed checking if index name is unique"	186051
Oracle Target	Issue	When the **Use direct path full load** option in the **Advanced** tab was disabled, the Full Load process would sometimes result in data loss when loading tables with numerous columns.	184903
SAP Application	Issue	The following types were sent as strings instead of numbers, resulting in changes not being applied to the target: 'a' - Decimal Floating Point Number, 16 Digits 'e' - Decimal Floating Point Number, 34 Digits	185878
IBM DB2 for LUW Source	Issue	The task would fail after capturing changes from a table with a LOB column, but without a Primary Key/Unique Index.	185588
Oracle Source - General	Issue	Modifying the "Expose number as" option in the endpoint settings, would propagate the change to all other Oracle source endpoints as well.	185328

Component	Type	Description	Reference
SAP Application Source	Issue	The task would fail while attempting to replicate a CDPOS table in Full Load.	185031
Hadoop Target	Issue	Describe of large tables with a primary key would fail due to insufficient buffer size allocated for the results.	184906
IBM DB2 for LUW Source	Issue	The task would sometimes fail while capturing a DML from a table with VALUE COMPRESSION set to the default value.	184441
File Source	Issue	The following error would be encountered when trying to capture changes from a file. `Partial results are valid but processing is incomplete`	184003
IBM DB2 for iSeries Source	Enhancement	Improved Change Processing performance.	183983
ODBC Source	Enhancement	Added support for unloading data with dBase ODBC driver in ODBC endpoint.	183920
Oracle Source-Attunity Log Reader	Issue	High memory usage was encountered when using the parallel ASM read option.	183910
Kafka Target	Issue	When replicating data from a mainframe source to Kafka target with logical avro datatypes, date fields with "0001-01-01" would be converted incorrectly.	183872
SAP Application	Issue	The search logic only used the first language found in a Business Group. Tables that used a different language would not be available for selection in Replicate.	183770
IBM DB2 for	Issue	Replicate would fail to parse CDC	183598

Component	Type	Description	Reference
iSeries Source		changes for a date column stored in *USA format.	
IBM DB2 for iSeries	Issue	Issues were encountered when trying to run CDC with the Journal set to *FILE (minimal logging), which isn't supported. Added to Help that changing the MINENTDTA Journal option from *FILE to *NONE is required for all CDC use cases.	183595
Oracle Binary Reader	Issue	Changes captured from IOT tables when one of the non-key columns contained a NULL value would return the following warning message: `Non key columns are ignored in Index operation for the table` This may sometimes lead to data inconsistency.	183529
PostgreSQL Source	Issue	An issue replicating from PostgreSQL was resolved by changing the SYS_PASSWORD_ LENGTH from 256 to 1024.	183418
Replicate General	Issue	Excessive memory consumption was encountered after upgrade, resulting in eventual task failure.	183309
Sybase ASE Source	Issue	Replication from Sybase ASE 12.5 would fail as Replicate would run a query for columns in the **syspartitions** table (**cdataptnname**, **indid**) that do not exist in that version. This query will now only be run from Sybase ASE 15 (which supports these columns).	183307
File Source	Issue	NULL values would not be copied correctly when reallocating the	182977

Component	Type	Description	Reference
		buffer used to read from the file. This resulted in the task stopping abnormally when rows exceeded 10965 characters.	
MySQL Target	Issue	When the task was set to Transactional apply mode, the target DATETIME value would be applied according to timezone instead of UTC.	182747
Metadata Manager	Issue	The **Optimize handling when LOB size is less than (KB)** feature would not work at table level when the table name or owner was changed in a transformation.	182746
Oracle Source - Attunity Log Reader	Issue	After running the Oracle Standby Failover procedure, the task would disappear from the Replicate Console.	182708 182787
IBM DB2 for iSeries	Issue	Changes to dates stored in *EUR format would be parsed incorrectly.	182592
Replicate General	Issue	After stopping a task and starting from timestamp, and then reloading a single table, the Change Table would not be recreated despite the **DROP and CREATE Change Table** option being selected.	182514
Oracle Source	Issue	Replication from Oracle source in a standby failover scenario would not work properly.	182420
Microsoft SQL Server	Issue	In rare situations, when the SP_ REPLDONE option was enabled, the task would fail to start with a query timeout when a huge transaction was performed on the source database. Added an option to increase the	182388

Component	Type	Description	Reference
		query timeout using an internal parameter.	
IBM DB2 for iSeries Source	Issue	Timestamps are adjusted according to the Replicate Server machine time. The issue was resolved by adding a **Time gap between Replicate and DB2 server** line to the log in seconds & microseconds.	182224
Microsoft Azure HDInsight Target	Issue	The following error was encountered during the task: `SYS,GENERAL_EXCEPTION,SSL validation failed.,CA file 'c:\temp\cacerts.pem' does not have read permission. Partial results are valid but processing is incomplete (apr status = 70008)`	181842 182003
MySQL Source	Issue	An "Unexplained Column Overflow" error would occur causing the task to stop. This resulted from the customer trying to start the task from timestamp after one of the source tables had been altered. This is not supported by Replicate - the Help was updated accordingly.	181814
IBM DB2 for LUW	Issue	During Change Processing, INSERTs and UPDATEs would not be applied to the target columns with codepage=0.	181336
Oracle Target	Issue	When applying changes to Oracle target, the task would crash under the following conditions: » Store Changes Processing is on	181319

Component	Type	Description	Reference
		» The **Limit LOB size to** value exceeds 64 KB » INSERT and UPDATE of the same record are executed in the same bulk	
Oracle Source - Attunity Log Reader Metadata Manager	Enhancement	Improved performance when processing many changes per second.	181314
Oracle Source - Attunity Log Reader	Issue	Attunity Log Reader Replicate would not capture Oracle Redo log events created using Oracle Global Transactions (dblink).	181225
Oracle Target	Issue	Replication to Oracle target would fail due to a missing code ICU page.	181146
Oracle Source	Issue	LOB lookup query on nullable columns would sometimes result in a full table scan.	181123
IBM for DB2 iSeries Source	Issue	Excessive memory consumption would occur during the task due to leftover Sorter swap files following a transaction rollback.	180829
PostgreSQL target	Issue	Replicate would fail to free connections that finished loading resulting in failure of loading subsequent tables due to insufficient connections.	180748
Sybase ASE Source Endpoint	Issue	Parsing a page number would return the wrong number when the captured event page number value was equal to or larger than 2 GB.	180724
Microsoft SQL Server Source	Issue	During CDC, tasks would fail with the following error: `Failure in accessing alternate backup folder`	180716

Component	Type	Description	Reference
Metadata Manager	Issue	Requests for the table definition of a table defined with dynamic partitions, would return a timeout from the server.	163947
IBM DB2 for LUW Source Endpoint	Issue	Connecting to a partitioned node other than "0" would crash the database instance.	163825
Oracle Source	Issue	Tasks would fail when a row with a LOB was already deleted during the LOB lookup, one or more of the source tables contains a BLOB column, and the following Replicate settings were defined: » The value in the **Limit LOB size to** field was less than 64 KB » The **Store Changes** replication option was enabled	163690
IBM DB2 for z/OS Source Endpoint	Enhancement	When an orphaned ECSA structure is detected, the associated UDF address space will be canceled via system command.	163641
PostrgreSQL Source	Issue	Resuming a CDC task may result in the first transaction not being applied to the target.	163611
Replicate General Table Statistics	Issue	The task would consume excessive memory due to over-allocation of memory in the Status Manager	163534
Oracle Source - Attunity Log Reader	Issue	The task would crash when a table with supplemental logging on all columns was updated vertically.	163501
MySQL Source	Issue	The `ResumeFetchForXRows` internal property was not working properly with VARCHAR	163472

Component	Type	Description	Reference
		primary key.	
Microsoft SQL Server source	Issue	When in BCP mode, the task would fail on TIMESTAMP column type when Unlimited LOB mode was enabled.	163259
IBM for DB2 iSeries	Issue	Excessive memory consumption would be encountered after processing numerous changes in the task.	163104
Amazon Redshift	Issue	When uploading a file to S3, a "Mandatory url is not present in manifest file" error would be encountered. The issue was resolved by adding a retry mechanism to Amazon S3 uploads.	162844
Teradata target	Issue	When TRACE logging was set for TARGET_APPLY, TPT messages would be written to the log causing a performance issue. The issue was resolved by writing TPT messages to the log only when the logging level is set to VERBOSE.	162467
MySQL Source	Issue	Date not-null 0000-00-00 from MYSQL source was replaced to 0000-01-01 for all targets instead of 0000-00-00.	161996
Oracle Source - Attunity Log Reader	Issue	Task would crash during the parsing of OLTP compressed tables with TDE tablespace encryption.	161729
Microsoft SQL Server Source	Issue	During CDC, due to erroneous parsing of compressed source records, null values would be applied to non-null target columns.	161383

Known Issues

The table below lists the known issues in this release.

Component	Description	Reference
IBM DB2 for z/OS Source	Tables are sometimes suspended when Replicate encounters a code 83.	185793
IBM DB2 for LUW Source	A restarted or resumed task may fail when capturing changes from a Db2 pureScale environment.	186487
ARC Source (IMS)	Parsing failures may occur sometimes due to invalid characters	185729
Oracle Target	When updating a record in a table with LOB columns fails to retrieve all of the LOB column values (possibly because the record had already been deleted), the BeforeImage row is not added to the Change Table.	186075
Oracle Source - Attunity Log Reader	The task may fail when the redo log that it was trying to access has not yet been written to the specified alternate folder.	182701
IBM DB2 for iSeries	Lock errors would sometimes be encountered during Full Load.	184019
Oracle Source - General	The "comments on columns" metadata query (used by the AEM Metadata module) may impact performance when replicating from an Oracle source.	185636
AIS (IMS)	AIS may sometimes not resume capturing changes after an extended period of inactivity in the source database.	180880