

Attunity Replicate 5.5 Release Notes - August 2017

Replicate 5.5 introduces several important features and enhancements including enhanced Change Processing support for IBM DB2 for z/OS and IBM DB2 for iSeries source endpoints, Change Data Partitioning, Load Ordering, and support for scheduling jobs according to UTC time or local server time.

Note In addition to these release notes, customers are also encouraged to read the release notes for all versions higher than their current version.

In these release notes:

- » [Migration and Upgrade](#)
- » [Enhanced Change Processing for IBM DB2 for z/OS and IBM DB2 for iSeries Sources](#)
- » [Change Data Partitioning](#)
- » [Load Ordering](#)
- » [Change Table Column Customization](#)
- » [Hadoop Changes and Enhancements](#)
- » [File Endpoint Enhancements](#)
- » [Support for Overriding Source Column Character Sets and Data Types](#)
- » [Attunity Avro Message Decoder SDK](#)
- » [SAP Application Source Changes and Enhancements](#)
- » [New Endpoints and Version Support](#)
- » [Rolling out Support for Attunity Replicate Connect \(ARC\)](#)
- » [Miscellaneous](#)
- » [Resolved Issues and Customer Requested Enhancements](#)
- » [Known Issues](#)

Migration and Upgrade

This section describes the issues that you may encounter when upgrading/migrating to Replicate 5.5.

.NET Framework 4.5.2

From Replicate 5.1.1, .NET Framework 4.5.2 or above needs to be installed on the Attunity Replicate machine.

User Permissions

In versions preceding Replicate 5.1.1, user permissions were defined in the **UserConfiguration.xml** file. From Replicate 5.1.1, user permissions are defined via the Replicate Web Console. To support this new functionality, when upgrading to 5.5 from a version prior to 5.1.1, the **UserConfiguration.xml** file is converted to a local repository file.

In most cases, the conversion of the **UserConfiguration.xml** file to a local repository file is straightforward.

However, although uncommon, the following case should be noted:

Users/Groups defined without a domain: If a user or group is defined in the **UserConfiguration.xml** without a domain name (e.g. `MyGroup` instead of `MyDomain\MyGroup`), the Replicate computer domain name will be automatically added to the user/group name (e.g. `AttRepDomain\MyGroup`).

If the Replicate computer does not belong to a domain, the computer name will be added instead of the domain name (e.g. `AttRepComputerName\MyGroup`).

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
```

Importing a JSON File from Previous Versions

Before importing a JSON file that was exported from a previous version of Replicate, you first need to edit the file and remove the **Scheduler** section (if set). Then, save the file and continue with the import process.

Kafka Target Endpoint

Replicate 5.2 introduces 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.

Oracle Source Endpoint

Customers upgrading from earlier versions 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.

Hadoop Target Endpoint

Importing a Hadoop target endpoint from earlier Replicate versions to Replicate 5.5 changes the Hive Access mode to ODBC.

If you were previously using WebHCat to access Hive, you should change it back to WebHCat after the import completes.

Teradata Database ODBC Driver

From Attunity Replicate 5.1.1, only Teradata ODBC driver 15.00 is supported. To prevent tasks from failing, customers with tasks that use the Teradata endpoint should upgrade their driver *before* upgrading to the new Replicate version.

Replicating System Names with the New IBM DB2 for iSeries Endpoint

The recently deprecated IBM DB2 for iSeries AIS-based endpoint replicated tables based on their system names (limited to 10 characters). The new native IBM DB2 for iSeries source endpoint replicates tables based on their SQL names (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).

Logging Levels

After upgrade, the logging levels are reset to the default level (Info).

File Target Endpoint

Post-processing commands that are imported into Replicate 5.5 from earlier Replicate versions will not run unless the **Post-process files** check box is selected (by default, it's not selected).

Note that this only applies to import/export and not to upgrade.

Upgrading from Version 4.0.9 to Version 5.5

Attunity Replicate 5.0 introduced significant changes aimed both at hardening the security of the product and ensuring its adherence to common industry standards. As such, several factors need to be taken into consideration when upgrading from Replicate 4.0.9 to Replicate 5.5.

This section highlights the main security changes and provides instructions on how to successfully upgrade from Attunity Replicate 4.0.9.

Security Related Changes

- » The technology of the Attunity Replicate web interface (Attunity Replicate Console) is now based on HTML5. Older versions used Adobe Flash which had a long history of security vulnerabilities.
- » In line with current industry security standards, the Attunity Replicate web server now enforces the use of HTTPS to protect against eavesdropping and data leakage. Using HTTPS requires a valid server certificate to be installed on the Attunity Replicate web server machine.
- » The default Attunity Replicate web interface URL was changed from:

`http://{host}/attunityreplicateconsole`

To:

`https://{host}/attunityreplicate`

This is due both to the new web interface technology and in preparation for an API to be released in the near future.

As a result of the above-mentioned changes, it is critical to refer to the installation chapter in the user guide and to make sure that all the required preparations are in place *before* beginning the upgrade.

Common Upgrade Issues

The following common issues may arise if an upgrade is done without preparation:

- » **Existing links to the Attunity Replicate Console do not work** - Since the new default URL is different, existing bookmarks and links will not work. This can be resolved by updating the links (recommended). Another way to resolve this issue is to change the URL back to the previous URL. This can be done by editing the `~Attunity\Replicate\Data\ServiceConfiguration.xml` file (default location) and changing the `AttunityReplicate` part of the `url=` attribute back to `AttunityReplicateConsole`.
- » **Opening the Attunity Replicate Console results in a 'connection reset' error** - HTTPS requires a valid HTTPS server certificate to be installed. If the installed certificate is not valid, the browser will not be able to connect to the Attunity Replicate web interface.
- » **On opening the Attunity Replicate Console, the browser issues a security warning about an untrusted web site** - When no HTTPS server certificate is found, Attunity Replicate generates and deploys a self-signed certificate. While this certificate is perfectly valid and secure, the web browser has no way of knowing this - hence the warning. This can be resolved by installing a valid HTTPS server certificate or by telling the browser it can trust that self-signed certificate. Note that the process for trusting certificates varies from browser to browser and your corporate security policy may prevent self-signed certificates from being installed (in which case you will need to purchase and install a third-party certificate).

Enhanced Change Processing for IBM DB2 for z/OS and IBM DB2 for iSeries Sources

Replicate 5.5 introduces design optimizations that improve performance and reduce the footprint when capturing changes from IBM DB2 for z/OS and IBM DB2 for iSeries.

Consequently, these endpoints can now be used without needing Attunity Integration Suite (AIS) to be installed on the source platforms. The ability to directly access these endpoints from Attunity Replicate reduces MIPS consumption, improves throughput, and reduces latency. These Change Processing optimizations complement the Full Load enhancements that were introduced in Replicate 5.2.

Change Data Partitioning

Replicate 5.5 introduces support for Change Data Partitioning.

Change Data Partitioning enables processing of Change Data from many tables in a consistent fashion. You can define the duration of partitions as well as the partitioning base time, thereby ensuring overall consistency of the partitioned data (i.e. no partial transactions).

The partitioned data is stored in the Replicate Change Tables. Information about the partitions is written to the new "attrep_cdc_partitions" Control Table in the target database. This information can then be used to identify partitioned data that needs to be further processed.

This feature is available in the Task Settings|**Store Changes** tab.

Note This feature is currently only supported with the Hadoop target endpoint.

Load Ordering

You can set the load order priority for each of the tables being replicated. Tables with the highest priority will be loaded first, followed by tables with lower priorities. This feature can be used, for example, for loading smaller tables before larger tables or for prioritizing load order based on data importance.

This facilitate this feature, a new **Load Order** column has been added to the table lists.

Change Table Column Customization

Users can now choose which header columns will be included in the Change Tables. By excluding unneeded header columns, customers can reduce the size of the Change Tables in the target database.

This feature is available in the Task Settings|**Store Changes** tab.

Hadoop Changes and Enhancements

The following enhancements have been made to the Hadoop endpoint:

- » The Hadoop target endpoint now supports sequence as a target storage format. When sequence format is selected, the data is loaded directly to the target (i.e. without running a Hive process).
- » Hadoop on Linux now supports Kerberos using a keytab file.
- » Change Data Partitioning. See [Change Data Partitioning](#) above.
- » **WebHCat Deprecation in Hadoop Target:** From Replicate 6.0, Attunity is planning to discontinue support for the WebHCat Hive access method in Hadoop Target. Customers currently using WebHCat to access Hive should start preparing to switch to an ODBC-based solution.

The ODBC driver that needs to be installed depends on your particular Hadoop distribution (Hortonworks, Cloudera, MapR, or Amazon EMR). For information about supported ODBC drivers, please refer to the Attunity Replicate Help.

File Endpoint Enhancements

The following enhancements have been made to the File endpoint:

- » **File Target:** The File Target post command now supports the following exit codes: Success (0), Recoverable Error (1), Table Error (2), and Failure (3 or above). Users can also use internal parameters to set exit codes with custom values. This is especially useful if the customer's application already uses the standard exit code values.
- » **File Source:** An **Advanced** tab has been added to the endpoint settings. The following options were moved from the **General** tab to the **Advanced** tab: **File preprocessing command**, **Check for changes every**, and **Change processing cleanup**.
- » Attunity Replicate data types are now included in the metadata file (DFM).

Note 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.

Support for Overriding Source Column Character Sets and Data Types

From Replicate 5.5, users will be able to override the character set and type of source columns.

This is useful in the following cases:

- » If a source character column is wrongly encoded. For example, if a source character column is described as encoded in CCSID X, but the data stored in that column is actually encoded in CCSID Y.

- » If a source column is defined as character type but the data stored in that column is binary or vice versa.

Note The following limitations apply:

- » Currently supported with IBM DB2 for iSeries and IBM DB2 for z/OS endpoints only.
- » NCLOB columns are not supported.

Note As Replicate never modifies the actual source data, any changes are applied during buffering.

Attunity Avro Message Decoder SDK

Attunity Replicate supports message streaming using both JSON and Avro formats. To ease the task of processing Avro messages, Attunity provides the Attunity Avro Message Decoder SDK, which contains the following:

- » Avro message deserializer
- » Working samples

SAP Application Source Changes and Enhancements

The following changes and enhancements have been made to the SAP Application source:

- » The SAP RFC NetWeaver client is now required to be installed on the Replicate Server. For instructions, see the *Attunity Replicate User Guide*.
- » The authorization object S_DEVELOP is no longer a required permission for the Attunity communication user in SAP.
- » The authorizations required for the dialog and communication users no longer require a strong basis role.
- » Changes to the structure of source SAP tables and changes to the SAP business groups are now handled automatically within Replicate.
The only exception is moving a table from one business group to another – in this case, the change needs to be updated manually in Replicate's **Table Selection** window.
- » The Monitor for Full Load now shows % complete based on the number of rows read from the SAP source as opposed to the number of rows transferred. This makes the estimation for complex SAP table structures more accurate.
- » Replicate now automatically reconnects to the SAP source in case of a connection failure.
- » A new **Connection Type** field was added to the endpoint settings. The new field enables users to choose either **Type A** (single server) or **Type B** (load balancing system) according to their SAP Application environment.

New Endpoints and Version Support

- » **Replicate for Linux:** Support for the Microsoft SQL Server endpoint has been added.
- » **Hadoop MapR:** Hadoop MapR 5.2 is now supported.
Note that SSL and Kerberos are currently not supported with MapR version 5.2.
- » **SAP Hana:** Support for SAP Hana as a target endpoint has been added.

Note The SAP Hana endpoint is currently in Preview mode. Endpoints or features in Preview mode are available for early preview, but are not generally available at the time of release.

Please contact us if you are interested in learning more about this endpoint.

Rolling out Support for Attunity Replicate Connect (ARC)

Starting from Replicate 6.0, AIS (Attunity Integration Suite) will be replaced with ARC (Attunity Replicate Connect).

As preparation for this change, the following endpoints appear in the UI with **(ARC)** instead of **(AIS)**.

- » DISAM (ARC)
- » IBM DB2 on iSeries (ARC)
- » IBM DB2 on z/OS (ARC)
- » IBM IMS (ARC)
- » RMS (ARC)
- » VSAM (ARC)

As stated above, the actual change in functionality is only due to be implemented in Replicate 6.0. Therefore, customers that are currently using AIS do not need to make any changes to their existing endpoints in Replicate 5.5. Customers should however be aware that in order to continue using these endpoints from Replicate 6.0, they will need to replace their AIS installation with ARC.

Miscellaneous

- » **Kafka Target endpoint:** Attunity Replicate data types were added to the metadata message.
- » **Deprecation of the Resume Full Load option:** The option to resume loading tables that were interrupted during Full Load has been deprecated. This option was accessible when the **Error** gauge was selected in the **Monitor|Full Load** tab.
- » The **Advanced Run Options** UI has been tweaked to remove redundancies and categorize the available options.

Resolved Issues and Customer Requested Enhancements

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

Component/Process	Type	Description	Ref #
Oracle Source - Binary Reader	Issue	Crash in parser of oracle binary reader would occur during an empty OLTP OVERWRITE event.	151488
Oracle Source - Binary Reader	Issue	INSERT of changed records that contained CR (Redo Chain) flags would not be applied to the target, despite being successfully applied to the Change Table.	149007
Oracle Source - Binary Reader	Issue	A task would fail when trying to replicate an OLTP compressed table, due to an unfamiliar Oracle redo log structure.	149800
Oracle Source - Binary Reader	Issue	A task containing compressed OLTP tables failed due to an overflow of an internal variable of columns.	149179
Oracle Source - Binary Reader	Issue	A task containing compressed OLTP tables using DIRECT INSERT failed to run.	148747
Oracle Source - Binary Reader	Issue	Fixed an issue with Oracle Source pertaining to updates that change the partition key.	148515
Oracle Source - Binary Reader	Issue	A task replicating partitioned tables would sometimes fail following operations on partitions.	147871
Oracle Source - Binary Reader	Issue	An error occurred when the redo log's physical block size was not a multiple of the logical block size.	147794
Oracle Source - Binary Reader	Issue	Oracle Binary Reader would fail when more than one archived redo destination ID was used.	147379
Oracle Source - Binary	Issue	Errors would be encountered	147116

Component/Process	Type	Description	Ref #
Reader		<p>when using Binary Reader to capture partition change operations.</p> <p>As Binary Reader does not support capturing partition change operations, the resolution was to start the task from a timestamp to clean the table partition list</p>	
Oracle Source - Binary Reader	Issue	The task would crash because of a previously unknown Oracle UPDATE event that Oracle creates in rare situations.	148925
Oracle Source Endpoint - Binary Reader	Issue	<p>Tasks copying archive directories would stop after two hours due to a timeout.</p> <p>The issue was resolved by adding a new internal parameter: <code>retryTimeoutInMinutes.</code> The default is 120 (2 hours). Setting "0" disables the timeout behavior and the task does not stop at all.</p>	N/A
Oracle Source - Binary Reader	Issue	An invalid timestamp error would occur on the target when the source tables resulted from SHRINK operations.	149979
Oracle Source - Binary Reader	Issue	<p>Due to an issue handling the V\$ARCHIVED_LOG entry, the task would stop with the following error (but could then be resumed):</p> <p>Redo archive file with the sequence 59980 does not exist.</p>	152182
Oracle Source - Binary Reader	Issue	When running multiple Oracle-to-Oracle CDC tasks (same source and target), the tasks would enter an infinite loop	153226

Component/Process	Type	Description	Ref #
		when reading the online REDO log.	
Oracle Source - Binary Reader	Issue	Due to an issue with the Oracle source endpoint, replicating from an Oracle source would result in an incorrect timestamp or incorrect numeric values on the target.	153463
Oracle Source - Binary Reader	Issue	Source latency would continue to increase when the database was idle	N/A
Oracle Source - Binary Reader	Issue	Tasks copying to archive directories would stop after two hours due to a timeout. The issue was resolved by adding the <code>retryTimeoutInMinutes</code> internal parameter. Setting the parameter to "0" prevents the operation from timing out.	N/A
Oracle Source - LogMinor	Issue	Oracle Binary Reader did not work when more than one archived redo destination ID was used.	151777
Oracle Source	Issue	"Test Connection" on Oracle source using Internal Parameter: "nlsCharsetId=874" would report an unclear error. The error was changed to: Cannot create Oracle NLS environment for charset 874 and ncharset 1000 OCI error.. SYS,GENERAL_EXCEPTION,Cannot create Oracle NLS environment for charset 874 and ncharset 1000 OCI error.,ORA-12715: invalid character set specified OCI error."	147814

Component/Process	Type	Description	Ref #
Oracle Source	Issue	High source latency would occur with Oracle source when ASM was used for storing Oracle redo logs.	148618
Oracle Source	Issue	chr(0) would be converted to single byte space instead of empty string on the target.	148507
Oracle Source	Issue	When the Oracle environment could not be established (e.g. due to an incorrect charset value), Replicate did not print the correct error messages.	147782
Oracle Source	Issue	Oracle source did not handle an internal flag correctly.	149251
Oracle Target	Issue	Replication would fail due to an Oracle bulk limitation of 330 columns per table.	147500
Oracle Target	Issue	An error would occur when dropping a column that is part of a primary key.	134893
Oracle Target	Issue	A double-byte space inserted into a column of type NVARCHAR2 was replicated as a regular space during full load. During Apply Changes, the task would stop abnormally.	147155
Oracle Target	Issue	Oracle bulk apply would sometimes crash during a Replicate upgrade when the net changes table already existed.	149147
Oracle Target	Issue	The char(0) of nchar and nvarchar2 data types would not be replicated correctly.	150255
Microsoft SQL Server Source Endpoint	Issue	The task would fail when reading a large transaction log file if the Replicate has file-level access to the backup log files option was enabled.	150370

Component/Process	Type	Description	Ref #
Microsoft SQL Server Source Endpoint	Issue	The user credentials for the Replicate has file-level access to the backup log files option would be refused.	151431
Microsoft SQL Server Source Endpoint	Issue	During change capture, the task would fail with the following error: A change record for table 'dbo.LOG_RAW_MSG' cannot be fully parsed.	151822
Oracle ASM	Issue	Performance issues were encountered when replicating from ASM.	148359
File Channel	Issue	When Full Load was aborted by the user, the File Channel remote task would not cancel the Full Load operation correctly. This resulted in missing events in the remote task's target database.	149221
File Channel	Issue	If the File Channel local task was stopped during Full Load with cached changes, the File Channel remote Task would start to apply changes on a table <i>during</i> the Full Load operation.	149221
Hadoop Source	Issue	Replicate would fail to capture files generated by Spark.	151186
Hadoop Target	Issue	An error would occur when configuring an ODBC connection using the Hortonworks ODBC driver.	149453
Hadoop Target	Issue	The task would fail as a result of dropping a column in a transformation.	148531
Hadoop Target	Issue	Setting high values for the "writeBufferSize" option would result in high memory	147932

Component/Process	Type	Description	Ref #
		consumption.	
Hadoop Target	Issue	An error would occur when trying to configure Huawei Hadoop as a target.	148609
Hadoop Target	Enhancement	Improved performance when using ODBC to access Hive.	149606
Hadoop Target	Enhancement	Added support for field delimiter values ranging from 200 to 377 octal (128 to 255 decimal).	153198
MySQL Source	Issue	Task would crash when parsing CDC strings with NLS characters.	149830
MySQL Source	Issue	The MySQL unsigned mediumint data type was incorrectly replicated to the target when the value exceeded the high boundary.	147961
MySQL Source	Issue	Full Load of a table with 4 billion rows would stall at 2 billion rows.	147486
MySQL Source Endpoint	Issue	<p>The MySQL <code>timestamp</code> data type would appear with an offset value on the target (e.g. 17:03:31 on the source would be 22:03 on the target). This is because MySQL converts <code>TIMESTAMP</code> values from the current time zone to UTC for storage, and back from UTC to the current time zone for retrieval.</p> <p>A workaround for this issue was provided using the following internal parameter:</p> <pre>serverTimezone</pre> <p>This should be set with the MySQL Server time zone.</p>	149504

Component/Process	Type	Description	Ref #
MySQL Source Endpoint	Issue	Stopping a task and then resuming it would sometimes result in missing UPDATES on the target.	150833
MySQL Target Endpoint	Issue	Bulk Apply would return "0 rows affected" when a key value was NULL.	145686
MySQL Target Endpoint	Issue	Memory leaks would occur when replicating data from a MySQL source to a File target on Linux.	152008
IBM DB2 for LUW Source	Issue	An unclear "offset X is larger than buffer length Y ignore" message would be displayed in the log. The message did not indicate that the source table was created with extended row size (and therefore not supported by Replicate). This was resolved by making the log message clearer.	148564 148566
IBM DB2 for LUW Source	Issue	Task would crash while parsing a variable string buffer.	148358
IBM DB2 for LUW Source	Issue	The prerequisite for the IBM DB2 for LUW driver was not correct in the help.	147700
IBM DB2 for LUW Source	Issue	The SELECT query to get the table metadata would wait on a lock and time-out when a table was being updated by a long transaction.	145352
IBM DB2 for LUW Source	Issue	The endpoint would incorrectly log a warning message when encountering a specific unrecognized event.	152004
IBM DB2 for LUW Source	Issue	Updating a source table to which columns were added would cause the task to fail.	152043
IBM DB2 for LUW Source	Issue	Capturing an IBM DB2 for LUW	153484

Component/Process	Type	Description	Ref #
		rare complex update would cause a crash.	
Microsoft Azure SQL Data Warehouse Target	Issue	An error would be generated when loading a string field with null value from a File Channel source endpoint.	147887
Microsoft Azure SQL Data Warehouse	Issue	Replicate's behavior when replicating source columns with empty values to Microsoft Azure SQL Data Warehouse was not clear. The Help was updated to make the behavior clear.	146718
Microsoft Azure SQL Data Warehouse Target	Issue	The <code>apply_exceptions</code> table would contain corrupt data due to incorrect use of memory pools.	148696
AIS Source	Issue	The value 24:00:00 in the DB2 for z/OS TIME data type would not be replicated.	147910
AIS Source	Issue	Events were not captured when Replicate AIS CDC started from 'now' and no CVS file existed.	147656
File Target	Issue	Archiving a file compressed with gzip (after processing) would fail.	147314
PostgreSQL Target	Issue	Task would crash during <code>target_lookup</code> on a PostgreSQL target.	147276
Teradata Target	Issue	A typo in the <code>td_data_encryption</code> internal parameter prevented communications from being encrypted between Replicate Server and Teradata.	150203
SAP Application Source	Issue	Table filtering would not work during Full Load.	149268
SAP Application Source Endpoint	Issue	An issue with the Oracle Binary Reader would cause tasks to	149943

Component/Process	Type	Description	Ref #
		fail.	
SAP Application Source Endpoint	Issue	Initial load of BSEG table would result in an error.	150777
SAP Application Source Endpoint	Issue	Capturing changes from a BSEG table would result in an error.	150978
SAP Application Source Endpoint	Issue	When replicating from a BSEG clustered table, Full Load would take a long time and Change Processing would not work at all.	151601
SAP Application Source	Enhancement	Replicating the Indexed Cluster Table STXL is now supported.	N/A
SAP Application Source	Issue	Due to an issue with handling non-consecutive column IDs, the task would fail after adding and removing columns to a table.	152150
ODBC With CDC Source	Issue	Removing and adding tables to a CDC task caused Replicate to crash.	149175
Sybase ASE Source	Issue	An error would be encountered when capturing data during a Change Processing operation.	149730
File Target	Issue	Resuming a task would fail when tables that were originally included in the replication were removed from the source.	150666
File Target	Issue	Task recovery would fail when tables that were originally included in the replication were removed from the source.	150877
File Target	Issue	The row count in the DFM file would not match the row count in the CSV file.	150844
Kafka Target	Issue	An unnecessary warning message about the table not having a primary key or a	148276

Component/Process	Type	Description	Ref #
		unique index would be generated.	
Kafka Target Endpoint	Issue/Feature Request	The Kafka target endpoint would fail to capture data from partial records. A limitation was added to the Help and a feature request was opened.	151233
Kafka Target Endpoint	Issue	Empty string values would be incorrectly applied as null values.	153657
UI - Transformations	Issue	Columns defined as NUMERIC (x,0) would be shown as NUMERIC(x,-1) after saving and then reopening the transformation.	148998
UI - Transformations	Issue	The UI would allow entering scale and precision for data types in Global transformations, even though the server does not support this.	149151
UI - Transformation	Issue	A global transformation rule converting from numeric to string would not work the first time it was defined.	147281
UI - Transformations	Issue	An error would be generated when trying to remove an expression from a transformation.	148388
UI - Transformations	Issue	The Output column order could not be reverted to the original table order. This was resolved by adding a "Restore default sorting" option to the UI.	148207
UI - Tables	Issue	When trying to edit a table, the UI would stop responding after removing a column using a transformation and then removing the column from the	148875

Component/Process	Type	Description	Ref #
		actual database.	
UI - AIS Source	Issue	The setEventsTimeOffsetSec property would not support negative values.	148863
UI - Endpoints	Issue	Endpoint names displayed in the console would be truncated when a task was selected.	149831
UI - Monitor	Issue	The UI would show the WARNING messages even hen clearing the warning message option and switching between tasks.	148143
UI Monitor - Full Load	By Design	<p>The "Estimated Finish Time" and "Time Remaining (Estimated)" values are not the same.</p> <p>Reason:</p> <p>The "Time Remaining (Estimated)" value is calculated by the combined transfer rate of all the records of the task, while the "Estimated Finish Time" is calculated per table. The discrepancy arises when the table transfer rate at the beginning of the task is very fast, but slows down towards the end of the task.</p> <p>In this situation, the "Time Remaining (Estimated)" value will be greater and less accurate than the "Estimated Finish Time" value.</p>	145848
UI Designer	Issue	<p>After changing the sorting order in the Table Settings > Transform > Output list, the original sorting order could not be restored.</p> <p>The issue was resolved by</p>	148530

Component/Process	Type	Description	Ref #
		adding a "Restore Default Sorting" option.	
UI Designer	Issue	Using the [Ctrl] key for Full Load and CDC Filtering would not work on a Mac. Now, on a Mac, you can perform these functions using the [Command] key.	152889
Sorter	Issue	After upgrading from Replicate 5.0.4 to 5.2, tasks would start to fail due to a conflict with sorter transaction versions.	150251
Sorter - Oracle to Kafka	Issue	Loading a table with many records while accumulating many cached events would result in unusually high memory consumption.	150366
Sorter and Logger	Issue	When recovering from a task crash, Replicate would sometimes start the recovery from the wrong place, resulting in missing data on the target.	150420
File Channel Sorter	Issue	When a task was stopped and restarted, transactions would be applied to the target at a much slower rate and the target would fall progressively further behind.	151334
Sorter	Issue	Replicate would not capture INSERTs in a remote task from Oracle to File Channel.	151361
Transformations	Issue	SQLite would generated an internal error when using a column name with a slash in a table transformation. The limitation was added to the help.	148543
Transformations	Issue	The outcome of transformation conflicts (Global Vs Table-	147804

Component/Process	Type	Description	Ref #
		specific) was not clear. A note was added to the help clarifying the matter.	
Transformations	Issue	An incorrect error message was displayed when failing to create table transformations.	147805
Transformations	Issue	Transformation for binary data types would not work.	147406
Installation	Issue	The installation would override the addons_def.json file.	N/A
Installation	Issue	Installation on Linux would fail when an existing group was used.	147811
Installation	Issue	Data directory would be created in the default location even if the server was defined to run with an alternative data folder.	N/A
Silent Installation	Issue	Silent install would not return a warning when prerequisites were missing. A note about the prerequisites was added to the Silent Installation section in the Help.	147945
Notifications	Issue	A error would occur when sending emails to an email server that does not support bare <LF>.	148004
Notifications	Issue	Notifications would not be modified after saving the changes.	147872
Upgrade	Issue	Upgrade to Replicate 5.1.2 failed in a Windows environment with FIPS encryption.	148078
Upgrade	Issue	Upgrading from 4.0.9 to 5.1.X or 5.2 would disable log roll over.	149469

Component/Process	Type	Description	Ref #
Installation	Issue	When the Replicate binaries and the Replicate "data" folder were installed on different drives, the shortcut to open the console would not work.	150935
Migration	Enhancement	Enabled migration path from a running server to a new server, without the need to reload the tasks.	147917
Logging	Enhancement	The Change Table names were added to the task log, enabling customers to delete them when a task ends.	147942
Logging	Enhancement	No messages were printed to the log when a task was in IDLE mode. This issue was fixed by printing a message to the log every 10 minutes when Replicate is in idle mode.	147068
Logging	Enhancement	The sequence values were not included in the TARGET_APPLY log.	N/A
Logging	Enhancement	Information enabling customers to measure the read time of each archive log was added to the SOURCE_CAPTURE log in DEBUG mode.	N/A
Documentation	Issue	The explanation for the Target table preparation "Do Nothing" option was not accurate. The Help was updated with a more accurate explanation.	149150
Documentation	Issue	A description of the metadata variables was missing from the Help.	148528
Documentation	Issue	Transformations would fail due to a known limitation. This	146287

Component/Process	Type	Description	Ref #
		issue was resolved by modifying the transformation limitations in the Help.	
Documentation	Issue	The "Stop the task after Full Load completes and cached changes have been applied" was working correctly, but appeared to be working incorrectly. A note was added to the Help explaining in more detail how this option is implemented.	147386
Documentation	Enhancement	A note was added to the Help clarifying that the Scheduler will run a task according the task's replication settings.	148994 146259
Documentation	Enhancement	Added a note to the description of the HQL scripts option in the Help, clarifying that the target storage format must be set to Text.	147874
Documentation	Enhancement	A note was added to the Help, advising customers to back up the Data folder before upgrade.	N/A
Documentation		Customer reported that a case number in the release notes was incorrect. It transpired that the case number was correct, but for a different customer (the same issue was reported by several customers). The customer's case number was also added.	151859
Documentation	Issue	Information for setting up the native IBM for DB2 for z/OS endpoint was missing.	149410
Documentation - Hadoop Target	Issue	Information about how to specify the cluster name when using a Knox Gateway was missing from the Help.	149902

Component/Process	Type	Description	Ref #
Task - Engine	Issue	When resuming a Full Load task from a timestamp, Replicate would try to run the task in Apply Changes mode, resulting in task failure.	150562 150823
Engine - Change Data Partitioning	Issue	Due to a possible issue with Control Tables being included in the batch partition, the task would fail with the following error: The transaction timestamp already exists in an earlier partition. The issue was resolved by including these tables in the partition and issuing an appropriate warning in the log.	N/A
Replicate Server on Linux	Issue	When using a custom data folder, the Replicate service could not be restarted after being stopped using "kill -9".	147983
Transactional Apply	Issue	It appeared that source changes were not being immediately reflected in the target. This was due to a failure to properly understand the impact of the Minimum number of changes per transaction option. The explanation of this option in the Help was expanded.	147526
Metadata Manager	Enhancement	Added support for reloading a table when the task is in a stopped state.	147409
Change Tables	Issue	The CT table column [header_ _]change_seq would be generated with a wrong timestamp value of 19700101 instead of the actual timestamp	148769

Component/Process	Type	Description	Ref #
Change Tables	Issue	When starting a task from timestamp, Replicate would drop and recreate Change Tables according to the policy defined in the task settings for "Store Changes." With this fix, Replicate will only create Change Tables if they do not exist when a task is started from timestamp.	149171
Change Tables	Issue	After resuming a task, the timestamp would be missing from the [header__] change_seq column.	151100
FTS Server	Issue	Stopping/reconnecting to the FTS server would sometimes cause a crash.	148303
Scheduler	Issue	The Scheduler date was not saved correctly.	146110
Run Options	Issue	Enabling the Metadata Only. Create Missing tables and then stop run option would have no effect. A Full Load operation would be performed and the task would not stop.	150700
Endpoint Shell	Issue	When retrieving tables, the endpoint shell would crash upon encountering source tables with a null owner.	150696

Known Issues

The table below lists the known issues for this release.

Component/Process	Type	Description	Ref #
Microsoft SQL Server	Issue	Support for AlwaysOn Availability Groups is not available when Replicate is installed on Linux.	N/A
File Transfer Service Logs	Issue	<p>Clicking the View Logs button in the File Transfer Service Logging Levels tab does not display the logs.</p> <p>Workaround:</p> <p>Browse to:</p> <p>C:\Program Files\Attunity\Replicate\data\repfts_defaultFts.log</p>	RPT-8980
Table Selection	Issue	<p>When clicking the Search button in the Select Tables window, the search continues indefinitely and no error is shown (although an appropriate error <i>is</i> written to the logs).</p> <p>This issue occurs when the user defined in the source endpoint settings does not have sufficient permissions to view the tables. For example, this issue will occur with the Microsoft SQL Server endpoint if the user does not have the db_owner permission.</p> <p>Note: In the event of such an occurrence, you can close the window by clicking the "x" in the top right corner.</p>	RPT-8246