



Syniti Data Replication (DBMoto) Release Notes for all 9.x Releases

Version	Date	Changes
9.8.0.19	7/22/21	<p>General/Common</p> <ul style="list-style-type: none">● Moved from .NET framework 4.6 to framework 4.8.● Removed "run as admin" requirement from Syniti DR client applications (Management Center, API, Service Monitor, Service Installer, LogViewer).● Moved client side application data files to C:\ProgramData.● Added support for Greenplum as a target.● Added support for CSV files as a source.● Integrated SDR Web Service for Microsoft IIS into the core product.● Added Apache Cassandra as a target and source(in Refresh mode).● Added an option to export Replication Browser information to Excel.● Integrated Refresh Partitions for use in setting up Refresh replications to achieve better performance.● Improved the metadata backup when the metadata.log abnormally increases in size.● Added a filter for Alerts on replication/group name for the replication Latency and Downtime.● Improved DTD property by calculating the absolute time difference with the database server.● Removed the Verifier Scheduler from the number of client connections● Introduced "ExecuteCommand" user permission to prevent execute statements to be executed if not allowed.● Renamed "Microsoft Azure Data Warehouse" to "Azure Synapse Analytics (formerly SQL DW)".● Added "Get Username from Session" Oracle connection property to get username from session if Syniti DR detects UNKNOWN user from log miner. Note: an additional query will be execute to get that information, impacting performance.● For Oracle connections, changed the default value for the Fetch Block Size to 1000000.● For Oracle, added support for "LONG" data type in bulk mirroring.● Improved support for sdo_geometry and other Oracle spatial data types.● For Microsoft Azure Data Warehouse (no primary keys), the Verifier can use record count only in this specific case.● Improved support for MS SQL Server geometry and geography data types.● Introduced AllowLoadLocalInfile property needed for bulk insert for MySQL connector 8.0.● Introduced unsigned int handling for MySQL.

- Replaced BinaryFormatter with a safer formatter.
 - Increased the log panel combobox size.
 - Improved Global Alert mail formatting and fixed an issue while editing email format when switching between text and html format.
 - Added an alert option when the Verifier scheduler reports differences
 - Introduced the option to call "Validate Replication" and "Rebuild Replication" from the external APIs.
 - For Metadata backup, added a function to restore from backup to main connection.
 - Added a column for Verification status in the Replication Monitor.
 - Updated the Replication Monitor grid to the latest available version.
 - In the Verifier Reconcile Data, added a message when completed even with errors.
 - In the Verifier, added support for Oracle default fetch size to improve performance.
 - In the Verifier, added support for comparing a subset of columns to improve performance. Primary keys are included by default.
 - In the Verifier, added max threshold on verification.
 - In the Verifier, a property to reconcile by operation type.
 - In the Verifier, added the option to reconcile only insert or delete or update operations.
 - In the Verifier, Source and Target connections are now mentioned in the details panel.
 - Fixed an issue with Identity columns for PostgreSQL.
 - Fixed an issue for replicating from PostgreSQL by using the Log Server Agent where the last ID was not available when creating a replication.
 - Fixed an error that occurred when storing metadata on Azure Synapse.
 - Fixed an issue found by the Security Code Scan plugin.
 - Fixed an error where Replication Validate did not remove the error flag from the replication.
 - Fixed an issue with Browse connection rule to allow user to run the Verifier.
 - Fixed an issue with LogViewer.exe app failing when using log4net for the Syniti DR log.
 - Fixed an issue where Oracle foreign keys not detected
 - Fixed an issue with Oracle where the default values defined on fields were not displayed in the Object Browser.
 - Fixed a null reference exception while editing Oracle connections when Log Server Agent was enabled.
 - Fixed an Obfuscation issue where VerifierProperties were not visible in the API.
- ### Management Center
- Fixed an issue where null exception raised when running the Management Center with no default metadata.
 - Fixed a regression bug where the Log Server Agent service did not start at the end of the wizard when the option to start it was checked.
 - Fixed a Mirroring filter compile error when using target fields not aligned with the source fields.
 - Fixed configuration issue where Teradata could not be added as source.
 - Fixed an issue with Snowflake where an error occurred while loading tables with mixed case.

- Fixed an issue in MS SQL Server target connections where an error occurred when setting up log reader.
- Fixed an issue with MS SQL Server Log Server Agent and Create Multiple Replication Wizard.
- Create table wizard is now creating a timestamp for a smalldatetime from MS SQL Server to Oracle
- Fixed the Verifier option "Verify Subset Of Columns" which not saved properly in the metadata.
- In the Verifier, fixed a problem with "Reconcile data error in replication". The Delete statement in some specific cases was not generated properly.
- Fixed an issue with IBM i (AS/400) where the function to retrieve unique indexes in tables was incorrect.
- Fixed an issue with Redshift when creating a connection with staging connection parameters.
- Fixed a regression bug in the Management Center Manage Licenses dialog where licenses could not be exported.
- Fixed a Management Center problem where the application did not scale correctly with non default DPI.
- Fixed a display issue with the Verifier Options dialog.
- Removed a warning message that was displayed when resetting the AWS S3 "Output Folder Archive" parameter.
- Fixed an issue where it was not possible to edit "Compare Subset of Columns option" from the Replication Properties Preferences.
- Fixed a Management Center error on disabling replication from right panel.
- Fixed an error with Reconcile for Multiple Verification: "...there are no cached data to reconcile..."
- Fixed an error with editing replications using Azure as target connection.
- Fixed an error that occurred when reopening connection to test with Azure Synapse Analytics (formerly SQL DW).
- Fixed an issue with table Aliases and catalog queries to load columns, data types, foreign keys and indexes successfully.
- Fixed an issue with the Custom Restore Wizard where the CommitMode and other options were not restored correctly.
- Fixed an issue with "Activate Log Performance" where a full path was not saved and the default "log" folder is saved.
- Fixed an exception opening the log file in some specific cases: "Non-negative number required."
- Fixed a communication error that occurred when loading big metadata backup files.

Replication Agent

- Added support for GUID data type from MS SQL Server Log Server Agent to MS SQL Server.
- Upgraded component dlls (ComponentPro SFTP) with latest version.
- For Kafka replications, included additional properties to support Kerberos authentication.
- Added support for JSON format as an Amazon S3 target.
- For Amazon S3 targets, added an option to use new file name format.

- Added a "Staging Refresh" feature that truncates the records in the staging table after the refresh is done.
- For SQL Server bulk mirroring, added the staging table to trace.
- Improved Delete Join statement for MySQL bulk mirroring.
- Added support for Teradata version 16.
- Added support for AutoInteger in bulk mirroring to MS SQL Server targets.
- Improved refresh replication mode by keeping data in the target table during the refresh.
- Added support for GUID data type with SQL Server Log Server Agent.
- Conducted stress test on some default parameters to find the best initial values. Increased default "Max Concurrent Connections" and "Max Concurrent Threads" from 20 to 50 and "Thread Execution Factor" from 10 to 100.
- Fixed an issue with using Snowflake bulk operations where the operation stopped after 5 blocks.
- Fixed an issue with S3 connections when the 'Output Folder Archive' was not set in the connection properties.
- Fixed a Replication Agent failure that occurred when a replication to Kafka or MongoDB failed to send data.
- Fixed an issue with Amazon S3 bucket or CSV target where no data was shown on record delete when "Add Transaction Info" = "NONE" (uncommon) was set.
- Fixed an issue where Snowflake did not properly handle case sensitive table with bulk mirroring.
- Fixed an issue with Snowflake where a bulk refresh error occurred when inserting empty string in a non nullable field.
- Fixed a bulk mirroring issue that occurred when creating a mirroring replication without primary keys.
- Fixed a duplicates issue with Azure Synapse.
- Fixed an error with MySQL in synchronization mode that occurred while inserting a new record when Override Identity Fields = True.
- Fixed an Informix DRDA bulk mirroring issue with Money data type.
- Fixed a case where PreExecuteCommands was not called in case of disconnection.
- Added a bulk mirroring recovery option for MS SQL Server during insert join. If it fails the Replication Agent switches to single insert just for that specific block.
- Fixed a MS SQL Server to Informix bulk mirroring issue with Boolean data type.
- For SAP HANA, fixed an issue with schemas in lower case in mirroring mode.
- Fixed an IBM i (AS/400) API Log Reading issue: Exceeded maximum number of SQL statements in package (1024).
- Fixed a PostgreSQL bulk mirroring bulk error: 42P01: relation "public.mara__staging_103" does not exist.
- Fixed a SQL Server error: LOB value reading error System.IndexOutOfRangeException: Index was outside the bounds of the array.
- Improved code to avoid Exception: Additional Info [Value: 123456xxx; Type: System.Int64; DestType: 2] ---> System.InvalidCastException: Specified cast is not valid.
- Fixed an "Index was outside the bounds of the array" error that occurred when reading LOB data from SQL Server when the column name contains the character '_'.

- Fixed a bulk mirroring bulk issue with Azure in tables without primary keys (uncommon).
- Fixed and created traces for bulk mirroring bulk with groups, error saving a record in the local cache (System.NullReferenceException).
- Fixed an issue with staging tables where sometimes in case of errors the LOG Message output had column names not aligned with values.
- Fixed a SQL Azure Data Warehouse datetime error that occurred during bulk mirroring. Datetime2 was not used properly in the staging table.
- Fixed an issue where the warning message "Perform INSERT if the UPDATE doesn't find the record" was not shown in synchronization.
- Fixed an issue where a replication does not start automatically when CommitID value is "0-1#0".
- Fixed a timestamp issue in IBM i (AS/400) when precision is 0 or 2 etc... not the default during mirroring.
- Fixed a SQL Server bulk mirroring issue "Cannot insert the value NULL into column". In this case the staging table was not using the right precision for a datetime2 data type and so the update join was failing to update some records.
- Fixed an issue with SAP HANA triggers that were created without quotes for column __ID.
- Added a check to prevent the Replication agent from stopping with "Object reference not set to an instance of an object" when DBMoto.Monitor.config is empty.
- Fixed an issue for target databases that do not support primary keys. If during refresh some very specific operations were performed the process could be terminated.
- For HANA triggers, fixed a special case when a fake gap occurred due to an unusual sequence number.
- For Oracle, fixed an issue that occurred when mirroring with a table without primary keys and with LOB fields.
- Fixed replication issues with non-standard mapping from time to datetime (System.Exception: The value "1753-01-01" is not valid for attribute...)
- Fixed an issue with PostgreSQL update statements using null parameters when switching to single insert mode after bulk mirroring failure.
- Fixed a MS SQL Server issue with bulk operations: "SqlDateTime overflow. Must be between 1/1/1753 12:00:00 AM and 12/31/9999 11:59:59 PM."

Log Reading

- Introduced the ability to schedule interruptions with the Log Server Agent.
- For Oracle Log Server Agent, implemented support for SDO_GEOMETRY.
- For Oracle triggers, implemented a trigger gap handling logic to prevent the fact that Oracle cannot use Read Uncommitted.
- Fixed a slow reading and timeout issue with IBM Db2 UDB Log Reader.
- Fixed an issue in IBM i(AS/400) log readers: When checking for MINENDTA supported values, combined values are now also checked.
- For Oracle triggers, fixed error "ORA-00905: keyword" that occurred when reading the last ID by datetime.
- Fixed an issue with MS SQL Server where creation of the distributor failed due to the generated password not meeting client's Windows security policy.

		<ul style="list-style-type: none"> ● Fixed a problem with Windows Authentication for the distributor installation on MS SQL Server. ● Fixed a login error with MS SQL Server Log Server Agent. ● Fixed an issue with the PostgreSQL Log Server Agent that occurred when inserting null values for bytea datatype with an empty bytea. ● Fixed a SQL Server Log Server Agent issue where all errors generated in the parsing process cause the Log Server Agent to stop. Parse problems are now skipped to continue to the next one.
9.7.4.5	6/21/21	<ul style="list-style-type: none"> ● [ALL] Addressed Log4Net vulnerability by upgrading the Log4Net assembly version installed with the product: "CVE-2018-1285 Unknown: log4net is vulnerable to XML external entity attacks. External DTDs are enabled by default and allow attackers to perform XXE attacks using malicious XML data and documents". The issue was reported here, and fixed in version 2.0.10 of Log4Net. <p>Management Center</p> <ul style="list-style-type: none"> ● Fixed an issue for MS SQL Server target connections when setting up the log reader. ● Fixed an issue related to SQL Server log settings: "Login failed for user "." when using Windows Authentication for the distributor installation. ● Fixed an issue causing failure to create distributor due to generated password not meeting client's Windows security policy. <p>Log Reading</p> <ul style="list-style-type: none"> ● Improved Oracle trigger algorithm by handling gap conditions due to Oracle not supporting Read Uncommitted. ● Fixed IBM Db2 UDB Log Reader issue causing slow reading and timeout error. <p>Replication Agent</p> <ul style="list-style-type: none"> ● Fixed an issue where Snowflake does not properly handle case sensitive tables with bulk mirroring ● Fixed an issue with Snowflake where a bulk refresh error occurred when inserting empty string in a non-nullable field ● Fixed an issue with bulk mirroring when creating a mirroring replication without primary keys ● Fixed a duplicates issue with Azure Synapse ● For SQL Server bulk mirroring, added a staging table to trace ● Fixed an issue with MySQL in synchronization mode where an error occurred on inserting a new record when Override Identity Fields = True. ● Fixed an issue with Informix DRDA bulk mirroring bulk and Money data type ● Added a bulk mirroring recovery option for SQL Server during insert join. If it fails, the Replication Agent switches to single insert just for that specific block.
9.7.3.3	4/9/21	<p>General/Common</p> <ul style="list-style-type: none"> ● Created deliverable images to run on Amazon EC2 via AWS Marketplace. ● DBMSserviceMonitor, removed error message when there is already an instance running. <p>Management Center</p> <ul style="list-style-type: none"> ● Fixed an issue with creating Redshift connections related to staging connection parameters. ● Create table wizard is now creating a timestamp for a smalldatetime from SQL Server to Oracle.

		<ul style="list-style-type: none"> ● Fixed an issue with the Manage Licenses dialog in the Management Center where it was not possible to export licenses. ● Fixed some graphics scaling issues with non-default DPI settings. <p>Replication Agent</p> <ul style="list-style-type: none"> ● Fixed an issue with IBM Db2 for i API log reading related to exceeding maximum number of SQL statements in package (1024).
9.7.2.28	1/26/21	<p>General/Common</p> <ul style="list-style-type: none"> ● Introduced a connection parameter for the maximum number of records to keep in memory while reading transactions. ● Added an alert when the Verifier scheduler reports differences. ● Improved message formatting for Replication Agent downtime alert. ● Included Ritmo/i version 7.1.0.3. ● Verified PostgreSQL partition table support. ● Improved browse connection rule to allow user to run the Verifier. ● Fixed the catalog query to get the primary key information from the IBM Db2 for i system table. ● Removed the Verifier Scheduler from the number of client connections. <p>Management Center</p> <ul style="list-style-type: none"> ● Added a column label for Transaction Latency values in Replication Monitor. ● Increased the size of the combo box in the Log panel. ● Added a warning in the Create Connection Wizard when choosing the wrong database type during Azure Synapse connection creation. ● Included the use of Extended Properties in the connection string for PostgreSQL LSA. ● Added a warning for the Oracle connection properties dialog where changing the password on the main connection does not change it on the log reader connection. ● Updated behavior of the "Override Identity Fields" option in Replication Properties. ● Set the "Transaction Latency Threshold" for each replication job, other than as global parameter. ● Source and Target connections are now mentioned in the Verifier details panel. ● In the Verifier, added support for Oracle default fetch size to improve performance. ● Redesigned the Verifier Options dialog to use tabs, fixing an issue where it was not displayed properly. ● Removed a warning message that occurred when resetting the AWS S3 "Output Folder Archive" parameter. ● Renamed Microsoft Azure Data Warehouse to Azure Synapse as DB Type in the Connection Properties dialog. ● Fixed an issue where last ID was not obtained when creating a replication from PostgreSQL using the Log Server Agent. ● Fixed an issue that occurred while retrieving PostgreSQL partitioned tables from the database. ● Fixed an issue where editing "Compare Subset of Columns option" was not possible in the Replication Properties Preferences.

- In the Replication Properties dialog, fixed an error that occurred when editing and saving multiple replications with an IBM Db2 source.
- Fixed an issue in PostgreSQL connection properties that occurred when changing the Max Pool Size and Min Pool Size (they should contain no spaces).
- Fixed an issue that occurred when setting Replication Agent trace filtered by replication: the trace was incorrectly enabled when adding or removing a replication from a group.
- Fixed an issue in Create Table Rule script where an exception occurred on creating a target table with target databases of type file (CSV, XML, JSON etc.)
- Fixed an issue that occurred when creating a target table from Oracle to SQL Server: timestamp to datetime2 would add an incorrect precision value.
- Fixed an issue with numeric data types (precision and scale undefined) in PostgreSQL.
- Fixed an error that occurred on disabling a replication from the Management Center right panel: Object reference not set to an instance of an object.
- Fixed a datatype issue when converting from long to varchar in the Create Table Wizard for Azure Synapse.
- Fixed an issue where a newly created Oracle connection was defined with Transactional Support = LogReader by default.
- Fixed an error with Reconcile for Multiple Verification: "...there are no cached data to reconcile..."
- In the Verifier, fixed a "Reconcile data error in replication" problem: the DELETE statement in some specific cases was not generated properly.
- Fixed an error that occurred in editing replication using Azure as target connection.
- Fixed an "object reference exception" error in the Connection Properties dialog when switching to a server with a different version than the one currently configured.
- Fixed an issue with table Aliases and catalog queries. Now columns, data types, foreign keys and indexes can be successfully loaded.
- Fixed an error related to a different script language in the metadata custom restore.
- Fixed an issue that occurred when setting up triggers: selecting existing master tables on schemas with no access permission would generate errors.
- Fixed the Custom Restore Wizard where the CommitMode and other options were not restored correctly.
- Fixed a datatype mapping issue in the Create Table wizard for Azure Synapse.
- Fixed an issue in the mirroring filter where the OR operator was not supported.
- Fixed an issue with the Start Page where links opened an Internet Explorer instead of the default browser.
- Fixed an object reference exception related to the Replication Monitor tab that occurred when loading the Management Center.
- Fixed an issue where in some cases the Management Center lost focus after clicking on a message box dialog.
- Fixed an issue with DTD where it was not saved in the connection properties dialog.
- In the Replication Monitor, removed msec portion from replication latency column.
- Adjusted the layout of the Enable Transactional Replication wizard.

Log Reading

- Added support for transactional replications using triggers with Progress OpenEdge as a source.

- For the IBM Db2 UDB log reader, added support for AutoXxx datatypes.
- For the Oracle log reader, fixed an issue that occurred in some cases when user ID is UNKNOWN.
- Introduced a new version of the Oracle Remote LSA with a simpler Java application that is easier to install and run on the database server.
- Added support for long passwords to the Log Server Agent for IBM Db2 for i.
- Fixed a problem where Log Server Agent config.ini files were not updated when mappings are changed via API.
- Fixed a problem that caused an error while parsing a SELECT LOB statement with a column value NULL.

Replication Agent

- Added support for replicating identity columns for Sybase ASE and SQL Anywhere databases.
- Added support for Progress OpenEdge database as a target in SDR replications.
- Implemented batch deletion of records for deleting all records from a table before running refresh for OpenEdge tables.
- Introduced refresh staging to improve refresh replication mode by keeping data in the target table during the refresh.
- Improved the DELETE JOIN statement for MySQL in bulk mirroring.
- Added support for Teradata version 16.
- Added support for AutoInteger in SQL Server bulk mirroring.
- Improved SAP HANA trigger-based replications.
- Conducted stress tests on some default parameters to find the best initial values. Increased default ""Max Concurrent Connections"" and ""Max Concurrent Threads"" from 20 to 50 and ""Thread Execution Factor"" from 10 to 100.
- Fixed a SQL Server LOB value reading error "System.IndexOutOfRangeException: Index was outside the bounds of the array."
- Fixed a Bulk Refresh issue that caused a failure in the recovery function when the bulk operation went into recovery: it could crash the Replication Agent.
- Fixed a bulk Refresh/Mirroring problem where, if an error was generated while recovering in single insert, the replication could not be disabled.
- Fixed an error (Index was outside the bounds of the array) reading LOB data from Microsoft SQL Server when the column name contains the character '_'.
- Fixed a bulk Mirroring issue with Microsoft Azure table without primary keys (uncommon.)
- Fixed a group bulk mirroring issue and added traces for a System.NullReferenceException error that occurred when saving a record in the local cache.
- Fixed an issue with staging tables where in case of errors the LOG Message output sometimes had column names not aligned with values.
- Fixed a SQL Azure Data Warehouse datetime error that occurred during bulk mirroring. Datetime2 was not used properly in the staging table.
- Fixed a replication scheduler issue where mirroring would start during a mirroring interruption period after rebooting the system running Syniti Data Replication.
- Fixed a regression bug recently introduced causing the Replication Agent crash after refresh.

		<ul style="list-style-type: none"> ● Fixed an issue with SAP ASE triggers where the identity gap when creating the master table is now explicitly defined (to avoid creating huge TID values in trigger tables). ● Fixed an issue in Oracle triggers to enforce order by where the transaction IDs were not ordered. ● Fixed a timestamp issue in IBM Db2 for i: when precision is 0 or 2 etc... not the default during mirroring. ● Fixed an issue where replication does not start automatically when CommitID value is "0-1#0". ● Fixed an issue with bulk mirroring for Microsoft SQL Server: "Cannot insert the value NULL into column" issue. The staging table was not using the right precision for a datetime2 data type and so the update join was failing to update some records causing an error. ● Fixed an issue where HANA triggers were created without quotes for column __ID. ● Fixed an issue in onExecuteError where, if the record was aborted, the error message was still written to the log. ● Fixed a PostgreSQL bulk mirroring error "42P01: relation "public.mara__staging_103" does not exist"
9.7.1.26	8/24/20	<p>General/Common</p> <ul style="list-style-type: none"> ● In the Verifier Options dialog, added a "Max Number of Comparisons" option to set a max threshold on verification. Default is 10,000. ● In the Verifier, enabled the option to compare a subset of columns to improve performance. Primary keys are included by default. ● For Oracle connections, changed the default value for the Fetch Block Size to 1000000. ● Added a "Get Username from Session" Oracle connection property to get a user name from a session if Syniti DR detects UNKNOWN user from the log miner. Note: an additional query will be executed to get the information, impacting performance. ● Added an optional connection property for Amazon S3 connections to specify a folder for the existing S3 Bucket. ● Renamed "Microsoft Azure Data Warehouse" to "Azure Synapse Analytics (formerly SQL DW)". ● Removed the "Override CCSID" property from the Ritmo/i connection. ● Fixed an issue where LogViewer.exe app crashed when using log4net for the Syniti DR log. ● Fixed an issue where Oracle foreign keys were not detected. ● Fixed an issue in the APILibrary where calling Enable/Disable/AsynchEnable/AsynchDisable on a grouped replication would modify the replication object reference. ● Fixed an obfuscation issue where VerifierProperties could not be seen from the API. <p>Management Center</p> <ul style="list-style-type: none"> ● In the Verifier, added the option to reconcile only insert or delete or update operations. ● For Microsoft Azure Data Warehouse where there are no primary keys, the Verifier using record count only is now allowed to run. ● Updated rebuild logic to show a warning for existing non-SDR triggers.

- Updated Sybase trigger logic (for Sybase 16.0) during Validate/Rebuild to check for existing Syniti DR triggers.
- Updated Oracle and Informix catalog queries to load table information correctly into the metadata when table name has an underscore.
- Fixed an issue that occurred when "Activate Log Performance" is checked where, when a full path for the log is specified but the default log folder is saved instead.
- Fixed an issue with the password in transactional setup for Oracle and SQL Server.
- Fixed a "Non-negative number required" exception that occurred when opening the log file in some specific cases.
- Fixed an error that occurred when reopening a connection to test with Azure Synapse Analytics.
- Fixed an issue in the mirroring filter where the Management Center would freeze if the filter could not be compiled.
- Fixed an issue in replication Validate/Rebuild where the logical Primary key in the source or target table would be removed.
- Fixed an issue in multiple selection from the Replication Monitor & Browser where an "Object reference not set to an instance of an object" was generated.

Log Reading

- For Log Server Agent, implemented log position file recovery.
- For Informix Log Server Agent, added an option not to turn off full-row logging.
- For Informix Log Server Agent, added TRUNCATE support.
- For Oracle Log Server Agent, implemented use of DefaultFetchSize property in mirroring.
- For IBM Db2 for i Log Server Agent, introduced support to read operations on table members.
- Fixed use of zero date time with MySQL Log Reader.
- Fixed a null char issue when pad mode is on with MySQL Log Server Agent.
- For MySQL Log Server Agent, fixed a null pointer exception when writing TRUNCATE command.
- Fixed an Oracle Log Server Agent remote error saving FTP setup dialog.
- Fixed an issue with MySQL log reading where column values containing space characters were not propagated correctly.

Replication Agent

- Added a check to avoid the Replication agent stopping with "Object reference not set to an instance of an object" if DBMoto.Monitor.config is empty.
- For MS SQL Server log server agent, added support for GUID data type.
- For Oracle targets, added support for "LONG" data type in bulk mirroring.
- Increased the default "Connect Timeout" value for MS SQL Server connections to 60 seconds.
- Modified mirroring to use FetchSize replication property.
- Fixed an issue for target databases that do not support primary keys. Certain specific conditions during refresh generated: "The process was terminated due to an unhandled exception."
- For SAP HANA triggers, fixed a special case where an unexpected gap was generated due to a weird sequence number.

		<ul style="list-style-type: none"> ● For Oracle, fixed an issue that occurred when mirroring a table without primary keys and with LOB fields. ● Fixed an error that occurred when enabling the performance log: "GUID-uniqueidentifier : Unable to cast object of type 'System.Guid' to type 'System.Byte[]'." ● Fixed replication issues with non-standard mapping from time to datetime (System.Exception: The value "1753-01-01" is not valid for attribute...) ● Fixed an issue with IBM Db2 for i source replications where the following error occurred: "ODBS Prepare Error: SYSPARTITIONSTAT in QSYS2 of type *FILE not found". ● Fixed an obfuscation problem with the script function GetJSONRecordInfo: "Unable to cast object of type 'fr' to type 'System.String'". ● Fixed an issue with PostgreSQL where, when switching to single insert mode after mirroring bulk failure, there was an issue with update statements using null parameters. ● Fixed an issue where newly added replications were not picked up by the Replication Agent if the client was closed before enabling them. ● Fixed a mapping conversion issue with MS Azure TIME datatype.
9.7.0.24	4/17/20	<p>General/Common</p> <ul style="list-style-type: none"> ● Added support for IAM role functionality for Amazon S3 in Redshift. ● Added support for IAM role functionality to connect to Amazon S3. ● Introduced Syniti Data Replication hosted help and added the option to use either the local or remote help. ● Introduced support for the CDATA ODBC driver for SAP RFC to connect to SAP HANA. ● Introduced the support for LONG VARCHAR datatype with Vertica. ● Introduced support for Kafka Confluent Cloud. ● Introduced support for the BOOLEAN datatype with PostgreSQL. ● Added support for Vertica version 9.3. ● Added an option to generate a JSON manifest file for Amazon S3 targets to be used by applications like Amazon QuickSight. ● Added support for Google BigQuery as a target. ● Added support for different regions with Amazon S3 buckets. ● Added support for staging connection for S3 settings with Redshift. ● Added support for Snowflake database as a target. ● Introduced a check when connecting to the Server Agent to verify if the client version installed matches the server version. ● Allowed the option to change the console language after installing Syniti Data Replication. ● Included the latest version of Ritmo/i 7.1.0.2. ● Removed a limitation on groups: now a group can contain replications having different target connections. ● Added support for Azure Data Lake Store Gen2 as a target. ● Added XML datatype support for IBM Db2 LUW when using the Log Reader ● Added XML datatype support for IBM Db2 LUW in refresh mode.

- Modified priority for Server Alert vs Replication Alert. Priority assigned to Replication Alerts. If the same alert is defined both on server and replication, the replication alert will take priority over the server alert.
- Enabled the ability to keep the customization of the SynitiDR Windows services (change startup type, login etc.) when upgrading to a new version.
- Added a license.txt file for the third-party components.
- Fixed incompatibility issues with metadata.log when upgrading from previous SDR versions.
- Fixed an issue with Custom Mapping where missing mapping records occurred in the metadata, which in turn would then trigger other errors.
- Fixed an issue with logging history related records(logs are saved to Database).
- Fixed a HanaDecimal issue with some target databases during refresh.
- Fixed an Informix DRDA issue that occurred when counting the number of connections.
- Fixed an error that occurred when reconciling data using the Verifier with synchronization replications.
- Fixed a regression problem with the Global Script where the AddLog global function was no longer visible from scripting.
- Fixed an issue with HANA staging table names; they are now generated using a portion of the metadata guid for uniqueness.
- Fixed an issue in the APILibrary where metadata got out of synch when creating multiple replications.
- Fixed an API library issue where IReplication.Enable returned a timeout error if the Replication Agent was running.
- Fixed incorrect downtime information in Data Replication Alerts.
- For Ritmo/i, fixed an error that occurred while reading source data during the data conversion function, with a table which included TIMESTAMP(3) columns.

Management Center

- Added the option to remove the table from the internal schema for XML File targets.
- Added the option to remove the table from the internal schema for CSV File targets.
- Virtual tables are now displayed in the Management Center using a different color and style.
- Added a new connection property for target connections to choose between DELETE and TRUNCATE during a refresh operation (if the target supports TRUNCATE).
- For IBM Db2 for i connections, introduced support for CCSID in create table operations.
- Introduced a check of the script language when a metadata is restored.
- Extended the connection object to include a staging connection when needed.
- Added a panel in the Replication Monitor to show summary values for Status and Transaction Latency.
- Added support for creating tables with Identity columns for the following targets: MS SQL Server, Oracle, IBM Db2 for i, MySQL, SAP HANA, MariaDB, IBM Informix, PostgreSQL, SAP ASE, SAP SQL Anywhere.
- Added a property to automatically refresh the metadata status from the Management Center.
- Fixed an issue where the Verifier options got mixed up if multiple Verifier tabs were open.

- Fixed an issue with the Verifier data grid where a default error dialog exception was displayed on closing the tab.
- Fixed an issue that occurred with the Oracle driver when reading decimal values: in some cases doubles were returned with a slightly different value and the verification failed.
- Fixed an issue with copy connection for Amazon Redshift where the staging S3 connection was being missed.
- Fixed an issue with Validate Replication for IBM Db2 for i LSA replications.
- Fixed an issue in the Management Center context menu, where the wrong right-click menu was displayed.
- Fixed an issue in the Replication Monitor where other windows did not respond correctly to a selected item change.
- Fixed an issue with saving Performance Logs to file(s).
- Fixed an issue in the Create Target Table wizard where the nullable primary key popup warning was not picking up the choice made by the user.
- Fixed an issue that occurred when installing the DBMOTOLIB on IBM i: if the user changed the name of the library, the new name was not used for the temporary lib.
- Fixed an issue in the Create Table wizard where the "Insert Field Before" and "Insert Field After" operations failed.
- Removed an incorrect message "Do you want to remove staging tables?" when scheduling a mirroring interruption.
- Fixed an issue in Create Primary/Secondary backup using PostgreSQL where there was an error saving metadata from memory.
- Fixed an issue in the Metadata connection properties where the primary backup database could not be changed.
- Fixed an issue with Triggers where validate/rebuild sometimes created a trigger with a wrong insert statement.
- Fixed an issue with showing Foreign Keys in the Object Browser for SAP SQL Anywhere version 9.
- Fixed an issue in the Expression Generator where the compare operator "greater than or equal" is now >= instead of =>.
- Fixed the "find" functionality for the Replication Monitor and Replication Browser.
- Fixed an issue in the Replication Agent Options dialog that occurred when trying to set up a database connection to save the log.
- Fixed an issue in the Replication Agent Options dialog that occurred upon changing the Trans. Latency Threshold value.
- Fixed an issue in the Create Multiple Replications wizard where BM Db2 for i tables with different journals were not handled properly.
- Added a warning in the Create Replication wizard for Trigger mode if trigger(s) already exist on the table.
- Fixed an issue that prevented changing the login to the Management Center.
- Fixed an issue with filtering functionality in the Log Viewer for languages other than En, It and Ja.
- Fixed a Verifier issue so that now when "Record count only" is selected, other options are disabled because they do not apply.

Log Reading

- Introduced support for read operations on table members with the IBM Db2 for i log reader.
- Introduced support for triggers with PostgreSQL.
- Added F-CE operation as a REORG operation on members for IBM Db2 for i log reader and Log Server Agent.
- For MySQL Log Server Agent, added an extra option `trace_table_list=<table1>,<table2>,...` to limit trace information to these specified tables.
- For Oracle Log Server Agent, added an extra option `trace_table_list=<table1>,<table2>,...` to limit trace information to these specified tables.
- For IBM Db2 for i Log Server Agent, introduced support for a case sensitive user password connection option and fixed install library procedure.
- Fixed a MySQL Log Server Agent issue with decimal columns on systems with "," as decimal separator.
- Fixed the error "Error adding a new replication to the log reader" when defining the Log Server Agent on MS SQL Server as source.
- Fixed an issue with Postgres Log Server Agent that occurred when comparing transaction ID values.
- Added a warning to the log file when MySQL Log Server Agent detects a table schema change.

Replication Agent

- Created AVRO serializer to be used with cloud storage platforms. Currently implemented with Google BigQuery.
- Added a counter in the performance log to show number of transferred bytes.
- Enhanced the existing scripting functionality to enable the user to access and update the target table name in CreateTableRule.
- Added an option to refresh the replications when a reorg event is detected.
- Fixed the error "XMLSchema:datetime is not declared" that occurred when replicating a source that contained a datetime field to XML.
- Enhanced the existing scripting functionality to enable the user to write functions in the Global Script that have access to Record Source and Target.
- Introduced a mirroring filter property to filter records in mirroring mode.
- Added support for TEXT datatype in Informix DRDA target for bulk mirroring.
- Fixed an IndexOutOfRangeException error that occurred when replicating to XML.
- Fixed the "FUNCTION MAX does not exist" error that occurred when using log to database with MySQL.
- Fixed an issue with the PostgreSQL Log Server Agent when using the test_decoding plugin: Error parsing column type "bytea" Unable to cast object of type 'System.String' to type 'System.Byte[]'.
- Fixed an issue that occurred when replicating from an IBM Db2 for i Time field to Oracle Date: System.Exception: The value "0001-01-01" is not valid for attribute "min_value".
- Fixed an issue with PostgreSQL as source in refresh: Can't convert .NET type System.String to PostgreSQL bytea.

		<ul style="list-style-type: none"> ● Fixed an issue when replicating to a File target where the number of processed records could be incorrect. ● Fixed an issue with the Replication Agent where it would stop working when it found an invalid dbmoto.monitor.config file. ● Fixed errors in the JSON syntax when replicating to JSON. ● Fixed an issue with MS SQL Server as source where replications got disabled when restarting SQL Server. ● Fixed an issue where refresh replications failed with Japanese system & Microsoft Access. ● Fixed an issue with bulk insert for BLOB datatype with MySQL as a target. ● Fixed the SAP HANA trigger syntax to drop a trigger if it exists. ● Fixed an issue where the IBM Db2 for i Log Reader API does not properly recognize the USERID in synchronization replications and this could create loops. ● Fixed an issue in mapping expressions where records were evaluated before for non primary key fields; instead, expressions should only evaluate record after values for non primary key fields. ● Fixed an issue where mapping expressions gave an error during mirroring when a Db2 for i table was journaled with AFTER image, though it worked fine if the table had BOTH images. ● Fixed an issue where errors were generated with PostgreSQL refresh bulk. ● Fixed an issue in refresh recovery: in some cases, when it recovered, an attempt was made to insert the last record again and a "Violation of PRIMARY KEY constraint" was generated. ● Fixed a Mirroring Bulk issue that occurred when the datatype was date time offset. ● Fixed an issue where in some cases a scheduled refresh did not start.
9.6.3.2	3/4/20	<p>General/Common</p> <ul style="list-style-type: none"> ● Introduced support for Kafka Confluent Cloud. ● Fixed an issue with the APILibrary where metadata got out of synch when creating multiple replications. <p>Log Reader</p> <ul style="list-style-type: none"> ● For Oracle Log Server Agent, added extra option to filter tables in a trace. ● Fixed an issue with MySQL log reader for decimal fields where records did not match when mirroring was performed. <p>Replication Agent</p> <ul style="list-style-type: none"> ● Fixed an issue that occurred when replicating an IBM Db2 for i Time field to Oracle Date. The following exception was generated: System.Exception: The value "0001-01-01" is not valid for attribute "min_value".
9.6.1.8	12/11/19	<p>General/Common</p> <ul style="list-style-type: none"> ● Added support for Azure Data Lake Store Gen2 as a target. ● Fixed a regression issue where the AddLog function was not visible from scripting. ● Fixed an issue where Syniti DR services were reset when there was an upgrade. ● Fixed an issue with the Ritmo/i provider for IBM i sources. when a table contained TIMESTAMP columns, the following error occurred "Error executing IResultset.GetValues while reading source data".

		<ul style="list-style-type: none"> ● Fixed an API library issue, where IReplication.Enable returned a timeout error if the Replication Agent was running. <p>Management Center</p> <ul style="list-style-type: none"> ● Fixed an issue with trigger behavior where validate/rebuild could sometimes create a trigger with the wrong insert statement. ● Fixed an issue with the Create Multiple Replications Wizard not handling Db2 for IBM i tables with different journals properly. ● Fixed an issue in the Replication Agent Options dialog that occurred when changing Transaction Latency Threshold and MailPassword value. ● Fixed an issue in the Replication Agent Options Dialog that occurred when trying to set up Syniti DR log using a database connection. ● Fixed the HELP/About dialog to show the correct version instead of 'unregistered version'. <p>Replication Agent</p> <ul style="list-style-type: none"> ● Changed logic to support SAP HANA triggers (it was previously not possible to use Read Uncommitted). ● For SAP HANA, added an identity column in Log Tables to handle uncommitted gaps. ● Fixed SAP HANA trigger syntax to drop a trigger if it exists. ● Fixed an issue where Mapping Expression did not work properly with Db2 for IBM/i table(s) journaled with *AFTER image. <p>Log Reader</p> <ul style="list-style-type: none"> ● Introduced support for a Case Sensitive UserPwd connection option and fixed the install library procedure for Db2 for IBM/i Log Server Agent. ● Fixed an issue where the IBM Db2 Log Reader did not handle user name with spaces correctly. ● Fixed an issue where the Oracle Remote Log Reader failed to write records when the SCN value was greater than 4GB.
9.6.0.50	9/12/19	<p>General/Common</p> <ul style="list-style-type: none"> ● Applied the new brand name "Syniti Data Replication" to the product. ● Changed the installation path and service names based on the new product name. ● Added support for MS Azure SQL Data Warehouse as target. ● Added support for MS Azure Data Lake Store as target ● Added support for Amazon DocumentDB as target. ● Added support for Amazon S3 as a target. ● Added support for MongoDB as a target. ● Added MariaDB to the database list. ● Added support for Kafka as target. ● Added support for XML database as target. ● Added Support for Amazon Redshift using the new Amazon ODBC driver ● Added Hadoop HDFS as a standard target. ● Introduced performance log to trace performance information while the Replication Agent is running. ● Improved Service Install utility to handle all Syniti DR services including the LSA.

- Implemented SAP Hana Time Travel feature.
- Implemented possibility to define Windows Authentication using AD groups.
- Improved security on database XML config files adding read-only signature to the files.
- Improved security on Syniti DR config files, adding signature to prevent manual editing.
- Fully supported Oracle ODBC driver
- Added bulk mirroring support for IBM Informix
- Added bulk insert for refresh to IBM Informix
- Upgraded product to .NET framework 4.6.
- Implemented bulk refresh using Array Binding with ODBC native APIs for TMaxsoft Tiberio.
- Added support for HIT ODBC/400 64-bit version.
- Added "Order By" sorting option when reading the journal on IBM Sb2 for i.
- Added support for replications status change inside the Verifier scheduler
- Implemented User Guide and APILibrary help in HTML5 format.
- Fixed Redshift issues with timestamp and timestampz data types aliases to specify the name in full (i.e: TIMESTAMP WITHOUT TIME ZONE)
- Fixed incompatibility issues with metadata.log from previous versions.
- Fixed an issue where temporary files were created but not deleted when opening the DBMoto.log from Management Center or the API.
- Modified GlobalScript.AddLog to parse the content of the message and link it to the replication or group object. Also added a new parameter to explicitly pass the replication or group name.
- Renamed some properties and changed values for FILE, JSON and HADOOP providers.
- For Hadoop, changed the default value for "Split Group Per Table" to true inside the config file. In this way the property will be saved properly inside the connection string.
- For Hadoop, a primary key on target table is no longer a requirement for mirroring.
- Fixed an issue that occurred when checking for a valid certificate for authentication.
- Fixed an issue in the alert system where the Replication Agent (Data Replicator) alert reported an outrageous number of hours.
- Modified setup and Management Center to uninstall LSA services & Restore service option when re-installing or restoring.
- Fixed an issue in the API Library, where, in some cases, the client application received a message saying that the metadata were unloaded.
- Fixed an issue in the API Library, where AutoFieldsMapping did not add mappings to the replication.
- Modified the value of the API ReplStatus group property when refreshing.
- Fixed a issue where the API SetInitialRefresh of a group did not work.
- Fixed an issue where the API ValuesAtGroupByOrder did not return the correct replication priority established in the group.

Management Center

- Improved the functionality of the Expression Generator dialog.
- Increased the maximum number of threads allowed by the Management Center.

- Improved the Replication Agent (Data Replicator) Settings dialog, showing proper warnings when a property value has been changed and the services need restarting.
- Added a property to skip the truncate before starting refresh.
- Added the isEngineUser property in the Users dialog.
- In the Management Center, added enforcement for the connection to use the proper database type.
- Made "Execute Initial Refresh" optional in the Replication Wizard.
- Enabled validate also for replications in refresh mode.
- In the Multiple/Export Replication wizards, added '<leave blank to use source table names>' in the Replication Name field.
- For Hadoop, the password/key_file is now obfuscated in the Connection Properties dialog.
- For Hadoop, improved documentation for the property "Path To Binary" that should include also the Hadoop command in the full path.
- In the Replication Agent (Data Replicator) Options dialog, Log tab, when you select the option to write the Log to Database, there is a new option:'Retention Time'. You can select one or both of (Keep Max and Retention Time) options. For example if you put 40 in the Retention Time, all records in the Log older than 40 hours will be removed whenever entries are made to the Log File.
- Updated the Verifier logic to include a "Refresh Filter Target" value when run.
- Fixed Japanese data display in the Verifier and SQL Query tab.
- Fixed an issue with the Multiple Replication wizard, where in Disconnected Mode the list of tables included all libraries.
- Fixed an unhandled exception while saving SQL Query.
- Fixed an issue in the Replication monitor where an unhelpful message was displayed.
- Fixed an issue that occurred while trying to open the Help in the Replication Options dialog.
- Fixed help strings in Replication Agent (Data Replicator) Options dialog.
- Fixed an issue that occurred while running the Enable Transactional Replication Wizard.
- Fixed an issue in the Options dialog where Preferences were not saved permanently.
- Fixed an issue in the Log Viewer where an exception occurred when opening the Verifier Schedule LOG file.
- Fixed an issue that occurred when using F5 on server or metadata. The node was not refreshed.
- Fixed an issue in the Replication Agent (Data Replicator) Options dialog, where some server configuration changes were reset when the dialog was edited.
- In the Metadata primary/secondary backup, fixed a null reference exception that occurred while configuring the backup connection.
- Fixed an issue where, in some cases, pressing the Show Replication History button returned an error.
- Fixed an issue in the Verifier where using primary keys only failed to reconcile the differences when there were multiple primary keys and the definition order was not matching the ordinal position.
- Fixed an issue where the maximum length of metadata identifiers was not checked.

- Fixed an issue in the Verifier where reconciling differences could generate the error: System.IndexOutOfRangeException: Index was outside the bounds of the array.
- Fixed an issue that occurred when opening history on a database: "The multi-part identifier '<dbname>.dbo.LogMessageID' could not be bound.
- Fixed an issue in Validate Replications where, in some cases, staging tables for bulk mirroring were not removed.
- Fixed an issue in the Verifier where reconcile data failed when using a refresh filter and choosing verify by PK only.
- Fixed an issue that occurred when opening the Dashboard: Error while opening log.
- Fixed an issue with Postgres Log server Agent where it was not possible to select an existing replication slot if using the test_decoding plugin.
- Fixed an issue in the Create Target Table wizard where the target table structure did not reflect the PK order of the source table.

Replication Agent (Formerly Data Replicator)

- Added support for replication of BLOB data when mirroring from IBM Db2 for i using the Log Reader API.
- Added support for a port number in Hadoop targets.
- Added support for Kerberos authentication in Hadoop targets.
- Added support for mapping from datetime to timespan with IBM Informix DRDA.
- Added support HanaDecimal in bulk mirroring.
- Added support for HanaDecimal in synchronization.
- Added support for MySQL Zero Date Time during bulk operations.
- Increased the default number of threads from 5 to 20.
- Added a new API for bulk operations with PostgreSQL driver version 4.x.
- Added support for a special case with Oracle bulk mirroring for tables with only primary keys.
- Added support for CreateTablePrefix and CreateTablePostfix when creating the staging tables for Oracle bulk mirroring.
- Added support for adding all record values if primary keys are missing, when available, in DELETE statements with Hadoop.
- Syniti DR now creates staging table for smalldatetime using datetime2.
- Added support for blanks in replication names with Hadoop.
- Added support for handling groups with Hadoop. By default, Syniti DR generates one file/table per folder. If a group is used, the folder name will be the group name, and files in the folder will be named according to table names.
- Added support for switching from bulk mirroring to single insert in case of errors with PostgreSQL targets.
- Implemented bulk mirroring for IBM Db2 LUW.
- Enhanced support for Event Logs. When selecting the option to save the Log file to an "Event Log", there are two options for the Log output destination: Syniti DR Log (default) and Windows Application Log. User can select the Windows Application log as the destination, but changes will be effective only after system reboot. All entries (Application level error and/or standard log entry) logged in the Windows Event Log will use the new source names.
- Fixed an issue where the Replication Agent (Data Replicator) did not properly recover the replication from a network error.

- Fixed an issue with Refresh recovery, where in some cases the PK index was incorrect, causing recovery to fail.
- Fixed an issue that occurred while compiling a replication script or mapping expression (System.IO.FileNotFoundException: Could not find file 'C:\Windows\TEMP\...')
- Fixed an authentication error that occurred when defining multiple users.
- Fixed an issue with the Replication Agent (Data Replicator) so that it no longer loops in situations like the following: "The record ... cannot be updated because according to a previous transaction it has been deleted."
- Fixed an issue in the Refresh scheduler where, despite scheduling every 2 week(s), the refresh was executed every week.
- Fixed a synchronization issue where a field was not found when accessing it from script.
- Fixed an issue with Informix DRDA bulk mirroring with a table that does not require update join.
- Fixed an issue with Informix "datetime year to second" bulk support.
- Fixed an Informix DRDA issue with special reserved words.
- Fixed an Informix DRDA issue with datetime datatypes in staging tables.
- Fixed Informix DRDA Delete Join syntax for bulk mirroring that was not working with more than one primary key.
- Fixed an Informix DRDA bulk mirroring issue that occurred with a datetime data type in staging table creation.
- Fixed an issue with Informix DRDA Serial data type in bulk mirroring mode.
- Fixed an issue with Informix DRDA bulk operations where now identity values (serial data type) are retained during a refresh (by default they were increased on the target every time.)
- Added a TRUNCATE statement to InformixDRDA.
- Fixed an issue with Oracle bulk refresh where a block with records in error caused the whole block to be reprocessed.
- Fixed an error that occurred when loading or processing the staging table in Oracle (Invalid identifier when adding a new column in the table.)
- Fixed an error on Oracle, "Input string was not in a correct format", that occurred when PDB not present.
- Fixed an issue with replication to Hadoop where the truncate operation did not remove files on the Hadoop side.
- Fixed an issue with Bulk Insert for PostgreSQL with .NET provider version 4.0.2.
- Fixed issues with MySQLDateTime data type.
- Fixed an issue in IBM Db2 LUW and z/OS where an error was generated when selecting the tables.
- Fixed an issue with IBM Db2 Bulk Mirroring (DISCONNECT in unit of work notallow STATE:25001 - CODE:-428.)
- Fixed an issue with the IBM Db2 for i log reader where DBMoto would write a warning to dbmoto.log for table reorganization only if the table name matched the member name.
- Fixed an issue with the IBM Db2 for i log reader where replication from ALIAS was not working correctly.

		<ul style="list-style-type: none"> ● Fixed an issue with RefreshFilterTarget where the filter was not applied for replications in a Group. ● Fixed an issue where an error occurred when replicating BLOB/CLOB with IBM Db2 for z/OS. <p>Log Reading</p> <ul style="list-style-type: none"> ● Introduced support for Postgres v.10 LSA replication, some system views have changed name. ● Improved LSA error messages regarding corrupted log position file. ● Introduced an option to support PostgreSQL on AWS RDS (Amazon) as a source for mirroring. ● Added log information with performance indicators to the LSA. ● For an Oracle Remote Log Reader, added a support database with ASM and additional local storage. ● Fixed an issue with Oracle standard Log Reader, where LogMiner was used without CONTINUOUS_MINE in versions equal or above 19c. ● Fixed an issue with Oracle LSA, where LogMiner was used without CONTINUOUS_MINE in versions equal or above 19c. ● Fixed an issue with the Oracle log reader where in some cases the user ID was missing. ● Fixed an issue with replications using Oracle 12c. It is now possible to enable the supplemental logging on a different container. ● Fixed an issue with MS SQL Server replications where an error was generated when parsing the SQL command from database log. ● Fixed an issue with the Informix LSA to continue and not stop when getting CDC_REC_ERROR record. ● Fixed an issue with MySQL LSA character columns that were incorrectly treated as UTF8. ● Fixed an issue with IBM Db2 for i LSA where an object reference exception occurred when configuring multiple replications. ● Fixed an issue with MySQL LSA where the port number was not stored in the config file and 3306 was always used.
9.5.5.15	3/19/19	<p>General/Common</p> <ul style="list-style-type: none"> ● For Microsoft Access, all unicode file names in a connection string are now handled. ● For Microsoft Access ODBC connections, added support for UNICODE characters. ● Fixed an issue with the API Library, where using AutoFieldsMapping did not add mappings automatically to a newly created replication. ● Fixed an issue in the API Mapping class that occurred when saving the mappings of a replication. ● Set APILibrary.dll to MSIL mode for all configuration builds. ● Fixed an issue with the ITable.SchemaRefresh function where it did not update LSA configuration files. ● Implemented SetDBMotoPath to redirect the assembly loading of the API code.

		<p>Management Center</p> <ul style="list-style-type: none"> ● Introduced a 'Create Index Postfix' property in the target connection to support extension to the index/primary key clause when creating the target tables. ● Added a "NO ALL" button to the dialog window for staging table removal. ● Fixed an issue in the Metadata Explorer, where clicking F5 on a server or metadata did not refresh the node. ● Fixed an issue in the Management Center, where the About dialog did not show the correct version details. ● Fixed a problem where a replication script was lost after multi-edit of replications. ● Fixed a Management Center locking issue that occurred when renaming metadata. ● Fixed an issue in the Verifier where reconciling using primary keys only failed to reconcile the differences when there were multiple primary keys and the definition order did not match the ordinal position. ● Fixed an "Index was outside the bounds of the array" error in the Verifier. <p>Log Reading</p> <ul style="list-style-type: none"> ● Improved MySQL LSA to allow only one cluster instance to run the LSA service at a time. ● MySQL LSA no longer parses newly added columns to source table. ● For Oracle LSA, introduced support for XMLTYPE columns. <p>Data Replicator</p> <ul style="list-style-type: none"> ● Fixed an issue that occurred when replicating with an IBM i RGZPFM (table reorganize) command: DBMoto used to write a warning to the DBMoto log only if the table name matched the member name. ● Added an option to use RID in place of RRN to improve performance. ● IBM i can now replicate BLOB data in mirroring using the Log Reader API ● Added an onMappingError global event to handle errors that occur during mapping. ● Managed a special case with Oracle bulk mirroring for tables with only primary keys. ● Fixed a problem in Replication alerts where log messages generated alerts also on replications that had no alerts defined. ● Oracle bulk mirroring now uses CreateTablePrefix and CreateTablePostfix for creating the staging tables.
9.5.4.3	11/14/18	<p>General/Common</p> <ul style="list-style-type: none"> ● For Hadoop, changed the default value for "Split Group Per Table" to true inside the config file in this way the property will be saved properly inside the connections string. ● For Hadoop, primary key on target table is no longer a requirement for mirroring. <p>Management Center</p> <ul style="list-style-type: none"> ● Fixed an issue with the Verifier in reconciling differences: in some very specific cases you could get the error "System.IndexOutOfRangeException: Index was outside the bounds of the array." ● Fixed an error that occurred when opening history on a database: "The multi-part identifier '<dbname>.dbo.LogMessageID' could not be bound." ● Fixed an issue where reconcile data failed by using refresh filter and verify by primary key only option.

		<ul style="list-style-type: none"> ● For Hadoop, the Password/Key_file value is now obfuscated in the Connection Properties dialog. ● For Hadoop, improved documentation for the property "Path To Binary" that should include also the hadoop command in the full path. <p>Log Reading</p> <ul style="list-style-type: none"> ● When a schema change occurs in the Log Server Agent, the replication property 'Stop On Schema Change' is used to determine the behavior of the replication. ● Fixed an issue with the Informix Log server Agent where it should but did not skip records that are rolled back with a DISCARD record. <p>Data Replicator</p> <ul style="list-style-type: none"> ● For Hadoop, added support for blanks in replication names. ● For Hadoop, by default one file/table per folder is generated. If groups are used, the folder name will be the group name and inside the file names will contain the table names. ● For Hadoop in mirroring mode, added all record values, when available, for a delete statement, if primary keys are missing. ● DBMoto now creates staging table for smalldatetime using datetime2.
9.5.2.6	10/17/18	<p>General/Common</p> <ul style="list-style-type: none"> ● Enhanced support for Hadoop by integrating it as a standard target ● Fixed an issue with metadata log growing indefinitely when using metadata on SQL Server. ● Fixed a conflict that occurred with changes applied by other users when scheduling a group for refresh. <p>Management Center</p> <ul style="list-style-type: none"> ● Fixed a thread synchronization issue when editing the metadata. ● In the Replication Wizard, fixed a null reference exception when a connection value was not set. ● Fixed an error that occurred when attempting to enable transactional replication with the Data Replicator running. ● Fixed an issue in Custom Create Table where changing data type did not work within the Replication and Multiple Replication Wizard. ● Fixed an issue where the replication script was lost after multiple editing of replications. ● Fixed an error that occurred when opening the Dashboard: "Error while opening log". ● Fixed an error that occurred when opening the Server Log with logs set on database: "Error while opening the log". <p>Log Reading</p> <ul style="list-style-type: none"> ● For SQL Server replications, fixed an error that occurred when parsing the SQL command from database log. ● Fixed an issue in validation of replications where the article not found and rebuild did not work in case of articles duplicated. ● Fixed an issue with the Oracle Log Server Agent where it would save only the first part of a LOB when the size was greater than 4MB.

		<ul style="list-style-type: none"> ● Fixed an issue in the MySQL Log Server Agent where the port number was not stored in the config file and 3306 was used always. ● Fixed an error with the Oracle Log Server Agent that occurred when creating a mirroring replication: "Input string was not in a correct format". <p>Data Replicator</p> <ul style="list-style-type: none"> ● Added an "Order By" sorting option when reading the journal on IBM DB2 for i. ● Improved data type handling with SQL Server replications: it used to handle a datetime2 as a datetime. ● Fixed an error that occurred while running a refresh bulk to PostgreSQL: "Keyword not supported: syncnotification". ● Fixed an "Index outside the bounds of the array" issue that occurred when using refresh recovery.
9.5.1.4	7/23/18	<p>General/Common</p> <ul style="list-style-type: none"> ● Introduced support for Tibero database as source in refresh mode and target. ● Introduced support for Oracle sdo_geometry and other spatial data types. ● Added support for SQL Server spatial geometry and geography data types. ● In the API library, added the ObjectID in the history message and allowed reading history without filter on replication. ● Fixed an issue in the DBMoto API where AsyncSetInitialRefresh used with AsyncEnable was not working. ● Fixed a Global Script issue where references were removed when clicking on Cancel or the X button to close the dialog. ● Disabled label edit/rename on the default metadata during execution to avoid an error on metadata getting out of synch. ● Fixed a Server Agent issue that occurred while attempting to cleanup logs on a database. ● Added support for CCSID 16684 in Ritmo/i ● Added HDFS target directory to Broker configuration. ● Added both logging and tracing to HBroker. ● Added ability to start HBroker as Application. ● Added configurable Hadoop binary path to DBMBroker. ● Created Hadoop service and added icon in the taskbar to start/stop service. ● Enabled Hadoop broker to move also mirroring files (*.mir). <p>Management Center</p> <ul style="list-style-type: none"> ● Added support for setting Transaction Read Point by a datetime for LSA sources ● In the Management Center, implemented a TAB key-press event in all grid controls ● Added trace details in the server agent, API and Management Center traces for operations that change metadata timestamp. ● Added record size information in the Table Properties dialog. ● Added a feature to the Verifier to verify a subset of records using a Pivot Column to filter records by last updated timestamp. ● Added the capability to allow reconciling records if the Verifier is run with Verify By Primary Key Only. ● Fixed an issue in the Verifier scheduler that occurred when changing the default metadata. ● Changed the DBMoto Verifier Scheduler service to Manual StartUp.

- Fixed an issue in the Group Properties dialog that occurred when changing any property and clicking OK or Apply.
- Fixed an issue in the remote Management Center where it would attempt to install the LSA on the remote server.
- Fixed an issue in trigger-based replications that occurred when switching the replication mode from mirroring to refresh for multiple replications: a message box was prompted for each of the replications.
- Fixed an error that occurred when setting up transactional replications on a File Provider connection. The message was "The specified folder for the provider does not exist".
- Fixed an error that occurred when reading from the history where log/history was set to database.
- Improved LSA dialog logic, specifically when LSA folder/INI file/LSA service were missing.
- Modified broker to catch failed command for copying to HDFS.

Log Reading

- Included a dbm_decoding.dll build for PostgreSQL ver 9.5, 9.6 and 10 (32bit and 64bit)
- For the Log Server Agent for IBM DB2 for i, introduced support for RRN.
- Log Server Agent for IBM DB2 for i now detects TRUNCATE and ALTER TABLE statement.
- Log Server Agent for MySQL will use all options specified in the connection's ExtendedProperties field when making a database connection.
- Fixed an issue with the Log Server Agent for Postgres where it passed the Command Timeout value to the INI file.
- Fixed an issue with the Log Server Agent for Postgres where it would replicate double byte characters as garbage.
- Fixed an issue with the Oracle log reader where there was an error parsing the SQL command from the database log: "grant SELECT on..".
- Fixed an error with the Log Server Agent for Oracle ("input string is not in a correct format") that occurred when creating a mirroring replication.
- Fixed an error with the Log Server Agent for PostgreSQL during mirroring for the "timestamp without time zone" datatype.
- Fixed an error with the Informix Log Reader "The third parameter passed to the UDR is not valid."
- Fixed an error with the Log Server Agent for SQL Server that occurred when a DML spans through multiple records in the log.

Data Replicator

- Introduced the FetchSize property in refresh replication properties when Oracle is source.
- Added TRUNCATE statement support for Oracle LSA, MySQL LSA and MySQL sources.
- For SQL Server, changed the type mapping from smalldatetime to datetime to have a more precise mapping conversion.
- Added support for the Guid data type in bulk mirroring.
- Added support for schema name for the staging tables for several mirroring bulk implementations.
- Modified the log messages to add the reason why a replication is being disabled.
- Implemented bulk insert on IBM DB2 databases.
- Fixed an issue in the Initial refresh of a single replication in a group (with Truncate Inverse = true), where it would run a truncate on all the other replications.

		<ul style="list-style-type: none"> ● Fixed an issue that would make the Data Replicator crash on a clean refresh of a group. ● Fixed a regression bug introduced in 9.5.0: initial refresh was not executed for the latest replication in the group. ● Fixed an issue for Informix timestamp values where IUD records are not kept but the values for COMMIT records are. The Replicator will use the COMMIT's timestamp value for the IUD records in the same transaction. ● Fixed an issue that occurred when replicating an XML datatype from SQL Server to SQL Server: "System.InvalidOperationException: SqlCommand.Prepare method requires all variable length parameters to have an explicitly set non-zero Size." ● Fixed an issue with MySQL triggers where the decimal part of a float was rounded in mirroring. ● Fixed an issue with bulk mirroring where a replication gets disabled if there are error conditions when loading the staging table. ● Fixed an issue related to when the Replicator reads a parsing error record written by the LSA. It will disable the replication if it detects a change in the source schema.
9.5.0.20	4/9/18	<p>General</p> <ul style="list-style-type: none"> ● Included Ritmo DB2 version 6.0.0. ● Added bulk mirroring for MySQL. ● Added support for the JSON datatype on all databases. ● Added support to manage schema changes on MySQL, SQL Server, Oracle, IBM Db2 for i, IBM DB2 UDB sources. ● Added continual messaging in the log settings to continue to write error messages in the log even during recovery. ● Added an FTP path property (SFTP) to be used during upload in certain FTP configurations ● Added an "Order By" Clause option in the replication property preferences. ● Implemented support for JSON file targets. ● Implemented 2 1/2 byte (HKSCS) support for IBM DB2 for i. ● Changed the DBMoto setup to install the unsigned version by default. ● In the Verifier Scheduler, added support for log to database. ● In the Verifier Scheduler, added support for persistent settings. ● In the Verifier, added support for some log values that could be retrieved (AS/400 RRN, Oracle ROWID...) ● When changing authentication mode, added verification for a specific authentication mode before switching to it. ● Updated the MD5 hash algorithm with SHA-256. ● Fixed an issue with alerts where the filter worked only on the top message string and not on the message stack. ● Fixed an issue in the Service Monitor that caused an error when stopping the Data Replicator on switching authentication type. ● Fixed an error that occurred while running the Verifier on MySQL with a timestamp column. ● Fixed an issue in the Verifier where it would give incorrect results with MySQL datetime.

- Fixed an issue with the Verifier Scheduler to resolve differences properly when multiple schedulers are defined.
- Fixed an issue where the "Reconcile Data" button was wrongly enabled during "Compare Primary Keys Only".
- Fixed a problem with the Verifier Scheduler where the log was not working.
- Fixed a problem with the x86 version of the DBMVerifierScheduler where a BadFormatException error occurred

Management Center

- Added a custom mapping function for global scripts to set Use Unmapped.
- Added menu item on a table to export the table structure (CSV, SQL, JSON format).
- Added an Initial Catalog property to the SQL Azure connection.
- Added a filter on group feature in Restore Replications and Export Replications wizards.
- Added a control to prevent the user from defining an unsupported synchronization replication for MySQL sources.
- Fixed an issue in the Oracle log reader where the Supplemental Log Database Level could not be set because the checkbox was disabled.
- Fixed an issue in the Validate Replications dialog where, with the Log Server Agent, the transactional setup installation was not verified correctly.
- Fixed an error in the Refresh Schema Information feature that occurred when updating a data type and the primary key settings at the same time.
- Fixed an error that occurred when loading the metadata after a reboot.
- Fixed an issue that occurred with high DPI screen settings where the Management Center thea wrong scaling, showing disproportional captions and labels.
- Fixed some issues that occurred when moving a replication from one log reader to another where the replication properties would still be set to the old log reader.
- Fixed the error "The DBMS selected is not supported by the current version of the product" that occurred when switching an IBM DB2 for i mirroring replication using the API log reader to refresh.
- Fixed a problem in Alerts. When the default metadata was changed, DBMoto would keep sending alerts for the metadata previously set as the default.
- Fixed the error 'ExecuteReader requires an open and available connection' in the Source Connection wizard, which was thrown when using the Test button after an error.
- Fixed a problem in the Enable Transactional Replication wizard with IBM DB2 for i, where it was not possible to leave the wizard step after removing the DBMOTOLIB path.
- Fixed a problem where the Create Table Custom Rule was not applied in the "Create Multiple Replications" wizard.
- Fixed a problem in the Export replications wizard where replications were added but the catalog/schema was missing.
- Fixed an issue with Metadata restore where an error occurred opening the replication field mapping.
- Fixed an issue in the History Viewer where elapsed time in some cases was displayed as '-: -: -, -'.

- Fixed an inconsistency that occurred when removing replications. In some cases (Hana/Netezza/Vectorwise) the user could not choose to keep transactional objects.
- Fixed a problem where sometimes the LSA replication Log File Name properties were displayed as blank in the Properties window.
- Fixed an issue where Verifier Scheduler settings were not saved properly using groups
- Fixed a problem in the History Viewer where the Session Time would sometimes cut off one day.
- Fixed an error in the SQL Execute Command panel that caused an invalid cast exception when loading datetime columns in some databases.
- Fixed a problem in the SQL Execute Command panel that occurred when setting number of rows to empty string.

Log Reading

- Implemented Log Server Agent support for IBM DB2 for i.
- Implemented PostgreSQL log reading for mirroring using Log Server Agent.
- Implemented Log Server Agent support for MS SQL Server.
- Introduced support for triggers for Sybase ASE v16.
- Added an option to re-create the LSA windows service after restoring a connection.
- For MS SQL Server log reader, added a check while reading records to verify if the Article ID matches the table name.
- Fixed an error that occurred when opening the log reader, DBMoto was attempting to delete .bin files that were locked.
- Fixed an issue with MySQL triggers, where an error was thrown if Allow Zero DateTime was set to true.
- Fixed an "Only one log reader agent can connect..." error with SQL Server log reader that occurred when dropping a publication.
- Fixed an issue with SQL Azure where a 'sp_helpsrvrolemember' error was generated when creating synchronization using triggers.
- Fixed an issue with the LSA binary log reader where transaction timestamps were incorrectly converted and sometimes not displayed in the replication status box.
- Fixed an issue where the Oracle LSA was hanging when a network cable or NIC was unplugged.
- Fixed an issue with the MySQL log reader where it would fail when the binary log file size was \geq 2GB.
- Fixed an issue where the Oracle LSA setup did not check for SUPPLEMENTAL_LOG_DATA_MIN = 'IMPLICIT' in addition to 'YES'.
- Fixed an error with the IBM DB2 for i LSA, where it was not working correctly for table names with more than 10 characters.
- Fixed a sporadic issue with the SQL Server LSA using boolean datatypes.
- Fixed an error with the IBM DB2 for i LSA where records in rollback to savepoint still got processed, resulting in duplicate key errors.
- Fixed a problem where the Oracle LSA sometimes did not get updated on the reset TID.
- Fixed an issue with the LSA where the service sometimes tried to delete wrong files under a specific condition based on the prefix value.
- Fixed an issue with the Oracle LRA/LSA where a NullReferenceException occurred when starting up the service.

		<ul style="list-style-type: none"> ● Fixed an issue where the Log Server Agent, enable/disable trace did not work while the agent was running. <p>Data Replicator</p> <ul style="list-style-type: none"> ● Added the global script function GetJSONRecordInfo to return a JSON formatted string from a record. ● Made Refresh Recovery persistent so that it could be used also when the DR is restarted. ● Fixed an inconsistency in GetReceiversInUse where receivers returned by this function were not aligned with the current content of the metadata, but only with the committed values in the database. ● Fixed an obfuscation issue with PostgreSQL bulk. ● Fixed an issue with PostgreSQL that generated a "Specified cast is not valid" error during bulk refresh with .net provider 3.x. ● Fixed an Oracle unique constraint violated error during mirroring bulk. ● Fixed an issue in bulk mirroring to MySQL with column names containing blanks. ● Fixed a problem where the merge operation for bulk mirroring looped when there was a duplicate record coming from the reading process. The warning is now logged. ● Fixed a problem with Postgres bulk refresh and 'float' data types. ● Fixed a problem where replications did not get notified when connection properties were changed from the Management Center and changes were not picked up by the Data Replicator when replications were enabled. ● Fixed an issue in the Data Replicator where, after stopping, some threads remained active. ● Fixed an issue in the Replication monitor where a request to disable a replication stayed in the Stopping status indefinitely. ● Fixed a problem with Netezza where numbers larger than the target connection could handle caused the replicator to hang on refresh. Problem records are now skipped. ● Fixed a problem managing the MySQL Boolean datatype when replicating from MySQL to MySQL. ● Fixed a problem in the Replication Monitor and in the history, where an incorrect total and failed number of records was shown when using Commitment Control. ● Fixed an issue with PostgreSQL Timestamp with time zone data type. ● Fixed an issue with the DBMoto Log for Oracle databases where the log is now able to recover and continue logging after connection issues.
9.0.9.7	1/19/18	<p>General/Common</p> <ul style="list-style-type: none"> ● Introduced support for SAP ODBC driver for Sybase ASE. ● Introduced support for IBM DB2 for i Alternative Catalogs (IASP), using the Catalog Name property in Ritmo/i and modifying the catalog query to specify in the syntax the alternative catalog ● Added an option to handle auto-incremental fields in the Data Replicator (currently only on Microsoft SQL Server and IBM DB2 for i as target) ● Added support for PostgreSQL driver version 3.0 ● Implemented remote Log Server Agent for Oracle ● Improved the alert "Reached Max Data Replicator downtime" by adding an option to set a value in hours and minutes.

		<ul style="list-style-type: none"> ● Applied a workaround for error "End of Stream encountered before parsing was completed" while loading metadata. A warning in the Agent Log is provided and the failing update is skipped ● Fixed an "Index out of bounds" error that occurred while running trigger replications using nested transactions. ● Added logging for MySQL Log Server Agent parsing error for a record: the error will be logged by the Replicator. ● Fixed an issue with MySQL bulk mirroring that occurred when column names contained blanks. <p>Management Center</p> <ul style="list-style-type: none"> ● Fixed an issue in the Management Center where, when applying mappings in synchronization, source fields mapped as expressions were not checked to set the Use Unmapped accordingly. ● Improved messaging in the Replication Properties dialog. When switching from Mirroring Bulk to Single Insert, the message asking whether to remove the staging tables would be prompted for every replication even if clicking the 'Yes All' button ● Fixed an issue with the Informix ODBC driver where refreshing table schema caused an InvalidCastException. <p>Data Replicator</p> <ul style="list-style-type: none"> ● Introduced a parameter in the onExecuteError event to put a replication into recovery mode ● Implemented bulk mirroring for PostgreSQL ● Fixed an issue where bulk mirroring was failing when the mirroring block was less than the default 10000 ● Fixed an issue with the merge operation for the bulk mirroring. It no longer loops when there is a duplicate record coming from the reading process. The warning is logged and the process continues. ● Fixed an issue in bulk mirroring to Oracle with column names containing blanks. ● Fixed an issue with Informix trigger_based replications to set column names to lower case if Identifiers case is set to Lower ● Fixed the error "LOB value reading error" that occurred while replicating LOB columns
9.0.8.40	11/6/17	<p>General/Common</p> <ul style="list-style-type: none"> ● Added support for bulk mirroring with MySQL ● Added support for bulk mirroring with MS Azure SQL and SQL Server databases ● Added support for bulk mirroring with Oracle ● Added the MirroringBulk property to the replication properties ● Added an option in Replication Properties to switch between different types of bulk refresh ● Included the latest Ritmo/i build ● Implemented HBroker for Hadoop replication in Windows environments ● Added a trace flag to write request action information (Enable/Disable or Set Pending/Set Executed status) to the Server Agent trace ● Added HANA as a source database for CDC/mirroring using triggers

- Fixed a problem in communication with the Management Center where a Log Server Agent replication's Log File Name was displayed as blank or never updated in the Properties window
- Fixed an issue in the Replication Manager thread where conflict occurred with changes applied by other users when moving replications in/out of a group
- Fixed an issue where the Server Agent service did not stop after new installation
- Fixed an issue in the Verifier Scheduler related to a deleted temp file after reconcile data
- Fixed an issue with the Oracle Log Server Agent where it incorrectly used the current SCN when starting up

Management Center

- Fixed an issue where a "Dock state is invalid" error occurred after an upgrade when launching the application
- Fixed a problem where the pending status of a group, when right clicking to expand the menu, was not aligned with the status shown in the Replication Monitor
- Fixed an issue with MySQL source connections where the Refresh Schema was not showing primary keys
- Fixed an issue in the History Viewer where, in some cases, log records were not filtered correctly based on the selected history item.
- Fixed an issue in the History Viewer where the Session Time would sometimes cut off one day
- Fixed an issue in the Replication Properties dialog where string properties were not saved properly when edited in multiple editing mode

Data Replicator

- Added support for schema names for the staging tables for SQL Server and Azure when performing bulk mirroring.
- Added an automatic rule to fix conversion errors from VARCHAR to NUMERIC.
- Added a feature to Oracle target connections to switch temporary to single inserts if the bulk mirroring is failing
- Added an ArrayBinding option for bulk mirroring with Oracle
- Implemented the Data Replicator trace on the fly, allowing users to change settings and replication filters while the Data Replicator is running.
- Added support for MySQL 'set' data type
- Modified mirroring and refresh behavior for SAP HANA so that DBMoto now generates internal files with unique names
- Added support for IBM DB2 UDB timestamp columns with fractional seconds other than 6
- Implemented Oracle bulk using the OracleBulkCopy class ("Native" option)
- Added a diagnostic setting in the trace for SQL Server log reader to print records from the transaction log.
- Added support for HKSCS implemented in Ritmo (DBC-12).
- Enhanced the Informix Log Server Agent to allow a single read of data changes for multiple databases of the same instance by specifying the same binlog folder and prefix for the different datasources
- Improved Oracle Log Server Agent handling of mismatched table columns. Warnings are now issued

- Handled insert on update in synchronization when the UPDATE comes from a DELETE+INSERT
- Fixed an issue where mirroring errors generated in the mapper did not update the transaction ID and caused transactions to be reprocessed
- Fixed an issue with bulk mirroring where Delete Join syntax was not working with more than one primary key
- Fixed a "Connection is already open" issue that could occur in specific cases where there were several exceptions
- Fixed a problem with Oracle bulk refresh by treating ORA-12592 and ORA-03111 as communication issues in order to put the replication in recovery mode
- Fixed a problem in synchronization by changing the DBMoto behavior in inserting records of duplicate primary keys
- Fixed a retroactive issue in the DBMoto log where it was getting written with newlines between messages
- Fixed an issue with the Replication Monitor where it would show an incorrect total and failed records when using CommitmentControl
- Fixed a problem managing the MySQL Boolean datatype when replicating from MySQL to MySQL
- Fixed an issue with replication scripts where an old version of onAfterExecute was not getting invoked
- Fixed an issue with synchronization where, in some cases, a replication or a group reset the target Transaction ID to null
- Fixed a "Cannot open CSV file" error in HANA
- Fixed an issue with the Informix trigger-based log reader where some errors were generated loading the master table when using Lower Case Trigger Identifiers
- Fixed an "Input string was not in a correct format" in the Oracle log reader which occurred when parsing CommitTID value
- Fixed an issue with in MS SQL Server trigger-based replications where a duplicated key error occurred when running an update
- Fixed a synchronization issue with MS SQL Server where in some cases the user ID of the transaction was incorrectly set to NULL
- Fixed an issue with MS SQL Server transactional replications where in some cases the transaction ID did not get updated
- Fixed an issue with Informix Log Server Agent to reduce frequency and stop treating log position file locking instances as errors
- Fixed an issue with Oracle Log Server Agent where, when a parsing error for a record is encountered, the error will now be logged by the Data Replicator. In the previous version, the instance was not logged
- Fixed an issue with MySQL Log Server Agent where the incorrect transaction timestamp was written
- Fixed an issue with MySQL Log Server Agent to make sure all database connections are closed after using them to read the log
- Fixed an issue that caused the Oracle Log Server Agent not to write errors to log position file under some conditions
- Fixed an issue where the Oracle Log Server Agent did not automatically pick up new tables when new replications were added

		<ul style="list-style-type: none"> ● Fixed a Log Server Agent issue where it attempted to delete the wrong files under a specific condition. ● Fixed an Oracle issue with a NullReferenceException when starting up. ● Fixed an issue where the LSA .ini file was not modified when refreshing a schema from the Management Center <p>Other</p> <ul style="list-style-type: none"> ● Fixed an issue in the APILibrary, where the SetInitialRefresh could become unresponsive.
9.0.7.6	4/21/17	<p>General/Common</p> <ul style="list-style-type: none"> ● Added Continual messaging in the log settings to continue to write error messages in the log even during recovery. ● For SAP HANA replications, added an FTP path property to be used during upload in certain FTP configurations. ● Added an "Order By" clause option in the Preferences tab of the Replication Properties dialog. ● Introduced support for SAP SQL Anywhere V17. ● Fixed an issue with the Verifier Scheduler to resolve differences properly when multiple schedulers are defined. ● Fixed an issue with the Oracle Log Reader that was causing LOB data to interrupt the Data Replicator. LOBs were also not replicated correctly. ● Changed Alerts logic to stop the Server Agent from becoming unresponsive when many log messages were queued waiting to be sent by the Data Replicator. <p>Management Center</p> <ul style="list-style-type: none"> ● Modified the Management Center trace to catch errors on reconciling data. ● Fixed an issue in the Management Center where replication lists may be emptied when clicking on some nodes in the treeview. ● Fixed an index out of range exception when opening the Log Viewer. Fixed some issues in the Replication Properties dialog: when editing multiple replications, some properties were not saved correctly. ● Fixed an issue where Verifier Scheduler settings were not saved properly when using groups. ● Fixed an issue in the Replication Properties dialog: when saving the last transaction ID, in some cases the CommitTID was not emptied out. ● Fixed an error that occurred when loading metadata on Oracle and the connection was down. The message reported that the metadata tables were missing. ● Fixed a problem in the Replication Properties dialog where string properties were not saved properly when edited in multiple editing mode. ● Fixed a problem that was causing slow performance while opening the DBMoto log on a database. <p>Data Replicator</p> <ul style="list-style-type: none"> ● Added SourceInternalTransaction and TargetInternalTransaction objects to script properties to expose the internal transaction object for custom code.

		<ul style="list-style-type: none"> ● Added a tracing option to write CommitTID and transaction information. Added support for NCHAR and NVARCHAR datatypes when replicating with SAP SQL Anywhere. ● Added support for new UPDATE FROM syntax when replicating with SAP HANA. ● Added support for an empty string when an import fails when replicating with SAP HANA. ● Fixed an issue with PostgreSQL and the Timestamp with time zone data type. ● Fixed a problem that occurred during replication where, if the CommitTID was set, in some cases when restarting a replication some transactions may have been skipped. ● Fixed an issue with the DBMoto Log on Oracle where it was not able to recover and start logging after connection issues. ● Fixed an error that occurred with Gupta SQLBase where the latest provider version expected parameters in a different format ● Fixed a null reference exception in a one-to-many refresh. ● Fixed an issue with HANA SFTP connections when a key exchange failed. ● Fixed an issue where the system blocked the Data Replicator an Verifier and would not allow them to run if FIPS was enabled.
9.0.5.6	1/26/17	<p>General/Common</p> <ul style="list-style-type: none"> ● For IBM DB2 for i Log Reader and Log Reader API, implemented a new stored procedure to accept to filter up to 300 files when reading the last ID. ● Implemented a new version of the IBM DB2 for i stored procedure able to handle up to 300 files with improved performance. ● Fixed a problem with the MySql Log Server: re-enabling a replication caused an System.NullReferenceException under some conditions. ● Fixed an inconsistency in the IBM DB2 for i Log Reader API where, in order to check library version and compatibility, the logic implemented for the standard Log Reader was wrongly used. ● Fixed a problem in synchronization with Oracle as a source, where, in some cases the user ID was read as UNKNOWN, causing the synchronization algorithm to fail. ● Fixed a problem that occurred during upgrade where the function that checks the SQL Server CE metadata would fail if no metadata had been created. ● Fixed an issue with the upgrade function where the AgentTimeout property was erroneously set to 10800. ● Fixed a problem in the verifier scheduler where it did not create the DBMotoVS_XXXX.LOG file. ● Fixed an issue in the Server Agent log where, in some cases, parts of the log file could be overwritten. <p>Management Center</p> <ul style="list-style-type: none"> ● Added a SQL Script tab in the Table Properties dialog to show the create statement for a source table. ● Fixed an issue in the History Viewer that occurred when the log was set to database and, in some cases, it did not show messages in order of timestamp. ● Fixed an issue with the SQL Query tab where, when switching user, the panel title would still show the initial login name. ● Fixed an issue that occurred While running the Management Center and the Data Replicator where, if the Data Replicator Settings dialog was edited from the

		<p>Management Center, an exception was generated saying that the operation was not completed because the metadata had been modified by another user.</p> <p>Data Replicator</p> <ul style="list-style-type: none"> ● Improved the Record_onBeforeExecute script function to allow an option for the Data Replicator to disable the replication before executing the current record. It is also possible to decide whether to skip the current record or leave the transaction ID at the current position. ● Improved refresh functionality when using SAP HANA by switching to single insert for a block when bulk fails during refresh. ● Fixed an issue in DBMoto.log where a wrong field name (obfuscation symbol) was shown in the log message when reporting a field mapping error. ● Fixed an issue with synchronization where in some cases a replication or a group reset the target Transaction ID to null. ● Fixed an issue where the log message did not identify the column that was having a problem in case of an "InvalidCastException: Specified cast is not valid." ● Fixed an error that occurred when replicating to an Oracle BLOB/CLOB/NCLOB: System.ArgumentOutOfRangeException: Specified argument was out of the range of valid values. ● Fixed an issue in transactional replications from an Oracle RAC where in some cases the replication could stay idle and not process existing transactions. ● Fixed an index out of bound array error in synchronization that occurred while performing the mapping operation. ● Fixed a problem that occurred when using MySQL bulk where the OnExecuteError event was not called in case of error.
9.0.3.7	10/5/16	<p>General/Common</p> <ul style="list-style-type: none"> ● Fixed a limitation that would prevent multiple users connected to a remote machine through Terminal Services from running the Management Center and Service Monitor in each session. ● Added a trace flag to write request action information (Enable/Disable or Set Pending/Set Executed status) to the Server Agent trace. ● Implemented a log file for the Server Agent. The agent log is always generated and contains information about connection, pending requests and errors from the Agent. ● Fixed a problem in ITable.RefreshCatalog where the operation did not send an update request to the Data Replicator and the replication was loaded with the old settings. ● For the MySQL Log Reader Agent, fixed an issue to support a binlog folder with spaces in the folder name. ● For the Oracle Log Reader Agent, fixed a problem in the setup to allow multiple datasources sharing folder and Windows service. ● Fixed an issue with the handling of the IBM DB2 for i catalog where Primary Key Extended and Foreign Key Extended information might not be retrieved correctly if Columns Extended was not checked. <p>Management Center</p> <ul style="list-style-type: none"> ● Fixed a problem in the Metadata Custom Restore and Export Replications wizards where some of the properties of the original replication were not exported. ● Fixed an issue in the History Viewer: when the log set to databasea, in some cases it did not show messages in order of timestamp.

		<p>Data Replicator</p> <ul style="list-style-type: none"> ● Fixed an issue that occurred during a bulk load on a SQL Server initial refresh of a huge table. If the bulk operation failed, when switching to single inserts, the connection was closed and the refresh failed. ● Fixed a problem that occurred when a group was in pending status but replications did not start refresh. ● Fixed a problem where pending refreshes in a group in some cases did not get triggered and the replication was left in Pending status. ● Fixed an inconsistency in synchronization where a DELETE and an UPDATE did not generate a conflict as the delete was always determined as a winner operation. ● Fixed an issue where the Verifier Scheduler could interfere with Data Replicator when sending request actions. In such cases, pending requests (Enable/Disable/Pending refresh) were not processed. ● Fixed a synchronization issue that occurred during the merge of operations, inside an interval, when the number of mappings was bigger than the number of table columns. ● Fixed an issue with MySQL Log Reader's parsing of the milliseconds part of DateTime columns on database version 5.7.
9.0.2.3	7/15/16	<p>General/Common</p> <ul style="list-style-type: none"> ● Added support for Oracle Version 10 as a source database with Log Server Agent replication mode. ● Implemented a new stored procedure for IBM DB2 for i sources to manage up to 300 files in transactional replications. ● The Metadata Custom Restore wizard can now be executed even if the Data Replicator is running. ● Fixed an issue with SAP Sybase SQL Anywhere where refresh replications failed because a wrong SQL statement was prepared. ● Refresh filters are now supported properly in the Verifier Scheduler. ● Added support for tables with a primary key logically defined in DBMOTO when using Oracle Log Server Agent transaction replication mode. ● Fixed an issue with the SQL Server log reader where in some cases the log reader did not detect the Log Reader Agent not running. ● Fixed an issue with the Oracle log reader: when reading from an old SCN, the complete list of files was loaded into the ADD_LOGFILE even if CONTINUOUS_MINE was set. ● Fixed an issue that was preventing the Management Center from reconnecting to the Server Agent when a connection was down. ● Fixed a problem that caused an error condition when an Oracle Log Server Agent source connection is an Oracle 12 instance installed without a container. ● Fixed condition that would make the Server Agent unresponsive, caused by errors while attempting to save a big metadata file. <p>Management Center</p> <ul style="list-style-type: none"> ● Introduced a feature to allow customization in the target table creation process. ● Added generation of a log history file for the DBMoto Verifier Scheduler. ● Added support for HanaDecimal data type in the DBMoto Verifier.

		<ul style="list-style-type: none"> ● Implemented an alternative method to retrieve catalog information from IBM i systems (DSPFFD and similar functions), for cases where the system columns table did not contain the necessary information. ● Fixed an object reference exception that occurred when setting up an alert on Reached Max Replication Latency. ● Fixed a problem where the Field Mapping could hang when creating automatic mappings for tables with several columns. ● Fixed a sporadic error causing the Management Center to hang when it was closed by pressing the X button. ● Fixed an issue where, when using Commitment Control, in some cases the number of records in the Replication Monitor and the history was incorrect. <p>Data Replicator</p> <ul style="list-style-type: none"> ● Implemented support for SQL Server Always On (Availability Groups.) ● Added support for use of SFTP when creating connections via the SAP HANA ODBC driver. ● Added support for the special character "/" in column names for SAP HANA connections. ● Fixed an issue that occurred during a bulk load on a SQL Server initial refresh of a huge table: if the bulk failed, when switching to single inserts, the connection was closed and the refresh failed. ● Fixed a problem that occurred when processing long bulk operations on Oracle: a disconnection error was thrown but the application was not able to recognize it as a disconnect, attempting to insert more records. ● Fixed an issue in the Data Replicator: if a replication was created Disabled while the replicator was running, it was never picked up for replication when it was enabled. ● Fixed a replication error that occurred when using the Commitment Control, causing the whole transaction to be rolled back. ● Fixed a regression problem that occurred when replicating LOBs with NULL values (Unable to cast object of type 'System.DBNull' to type 'System.String'). ● Fixed a Commitment Control issue when updating several records at a time. ● Records are now handled properly during the merging of operations inside a mirroring cycle when records are inserted and then updated (including the primary key) and then inserted again with the original primary key. This applies to the following target databases that use staging tables: IBM PureData/Netezza, HP Vertica, SAP HANA, Teradata, Vectorwise, SAP SybaseIQ.
9.0.0.10	3/15/16	<p>General/Common</p> <ul style="list-style-type: none"> ● Added Log Server Agent for reading MySQL binary logs. ● Added support for the Teradata database as a source database using refresh mode, and a target database using refresh and mirroring modes. ● Added Log Server Agent for reading Oracle logs using Log Miner. ● Implemented Oracle transactional support using triggers. ● Introduced a new logreader for IBM DB2 for i that is based on IBM APIs to support mirroring BLOB and CLOB data. ● Implemented the DBMoto Verifier Scheduler service, allowing users to schedule replication verifications and reconcile differences between source and target databases, if any.

- Added support for replication to Hadoop.
- Added support for SQL SERVER 2012 SP2 (CU7), SQL Server 2014 RTM (CU10) and SQL Server 2014 SP1 (CU3).
- Configures a Java-based version of the DBMoto Log Server to read redo/archive log files directly on the local system where it is running.
- Fixed a problem in the Oracle log reader for Oracle version 9.0.1 where setting up supplemental logging would generate an "ORA-00904: invalid column name" error.
- Fixed a problem in the Oracle log reader where some transactions at the end of a mirroring session could be reprocessed at the beginning of the next one.
- Fixed an issue with triggers for Oracle where new triggers could not be created if the table had already triggers defined.
- Improved the Oracle log reader on Oracle 9i: when the Data Replicator attempted to get the last ID at the end of a refresh, it would not get the very last items causing some transactions to be processed with a delay of several mirroring sessions.
- Fixed a problem with the Oracle (standard) Log Reader: a duplicate key error occurred in synchronization where the username read from the transaction log was UNKNOWN and caused transactions to be re-processed back to the source.
- Improved Oracle trigger-based replication; implemented and tested trigger-based replication on Oracle 10g.
- Fixed a Ritmo/i problem where a error occurred with Bulk Insert on IBM system i with GRAPHIC datatype.
- Changed behavior of the Alerts mechanism: email is no longer sent for disabled replications.
- Fixed a setup problem: "This implementation is not part of the Windows Platform FIPS validated cryptographic algorithm."
- Fixed a problem which caused a SQL parse error while reading a transaction from the DBMoto Log Reader for Oracle.
- Fixed an unhandled exception that occurred when opening the Toolbox from the Management Center main menu.

Management Center

- Introduced a new look and feel with new icons and a new graphical layout for the main window.
- Improved wizard windows, enlarging their default size and making them resizable.
- Implemented a warning message and an error for metadata that fails to commit changes to database.
- Implemented a way to schedule refresh replications to recur at the end of the month (e.g., last day of the month or last 3 days of the month)
- Added an Apply button to the transactional setup dialogs.
- Added a new IME Mode property in the Management Center Options dialog to set the input mode for the script editor dialogs.
- Added 'No To All' button to the prompt message when removing transactional replications.
- Changed DBMoto Log Reader captions to use more friendly names.
Fixed a problem that occurred when opening the Log viewer and the DBMoto log was set to database: the record info (field values) were not displayed.
- Fixed a generic error that occurred in GDI+ while running the SQL Command dialog and opening a table with CLOB data in it.

- Fixed an issue in the Replication Monitor, when the record to update or delete was not found on the target, the record counter showed one processed record and one failed record.
- Fixed a problem in the Management Center where it would sporadically hang and had to be shut down through the task manager.
- Fixed an error that occurred when parsing the query to set the DBMoto Log to a database.

Data Replicator

- Added a flag on Record_onBeforeMapping, Record_onAfterMapping, Record_onBeforeExecute and Record_onAfterExecute to allow the user to disable the replication.
- Added a parameter to the Record_onExecuteError function to allow disabling the replication from a script.
- Implemented support for commitment control where single commands are committed or rolled back following the source database committed transactions.
- Implemented the replication of LOBs via files instead of keeping their values in memory.
- Added support for LOB datatype in mirroring
- Fixed an error that occurred when setting the DBMoto to log to database: one of the log table columns was under-dimensioned and could cause an error when writing a new record to the log.
- Fixed a problem "Invalid Day value -1" that occurred when running a replication with a monthly schedule set to day 31 of the month
- Fixed an issue that caused a truncation statement to fail when attempting to clear an Oracle target table during refresh
- Fixed a MySQL fractional-time truncation issue in mirroring mode
- Fixed an issue where a group was disabled when the target table was an ALIAS and a connection error occurred during replication.