

DBMoto Release Notes for all 9.x Releases

Version	Date	Changes
		<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 AsynchSetInitialRefresh used with AsynchEnable 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 postgresSQL 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.
- 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.

		<ul style="list-style-type: none"> • 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.
- 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. • 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

		<ul style="list-style-type: none"> • 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
- Fixed a Log Server Agent issue where it attempted to delete the wrong files under a specific condition.

		<ul style="list-style-type: none"> 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.

		<ul style="list-style-type: none"> 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 Management Center, an exception was generated saying that the operation was not completed because the metadata had been modified by another user. <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 database, 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.

		<ul style="list-style-type: none"> • 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. • 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.