DBMoto Release Notes Version 7.x

Version	Date	Changes
7.2.1.2	**01/12/13**	• Added the property SingleTableCommit for a Netezza target connection. **Management Center** • Added support to the Verifier to run 'select count' only on source and target tables for the verification of one or more replications; • Fixed a problem in the Verifier where, in some specific situations when the "Verify Primary Key Only" option was specified, the Verifier did not create the select statement properly; • Added libraries containing stored procedures for DB2 UDB log based replications. **Data Replicator** • Solved a performance issue in synchronization; • Improved performance when the target is Vectorwise; • Fixed a problem where a refresh to a Vectorwise target could fail if the ADO.Net was not registered in the GAC; • Fixed a problem that occurred when mirroring to a Netezza target. An error loading or processing the staging table could occur if the target table had non nullable fields. • Fixed an issue that could cause DBMoto to read Informix log records incorrectly, following a missing before-update or after-update record; • Fixed a problem where columns that were used in mapped expressions, and not directly, could incorrectly be excluded from data capture because the column information wasincorrectly specified in the config file; • Added libraries containing stored procedures for DB2 UDB log based replications.
7.2.0.3	**10/26/12**	• Added support for Actian Vectorwise v2.5 and above as a target database, using the Ingres .NET Data Provider 2.1.1010 or later. **Data Replicator** • Improved performance when getting the last transaction id on SQL Server log based replications; • Added DB2UDB Log Reader support for compressed tables; • Fixed an error that occurred if processing a change from an Oracle source and the statement contained the UNISTR function; • Fixed an IndexOutOfRange exception that could occur during synchronization when an UPDATE changed primary key values;

		• Fixed an issue with the way DB2UDB Log Reader handles transaction ID numbers, so that DBMOTO does not cause error messages to be written into the DB2 diagnostic log file. If you are using log-based replications with IBM DB2 UDB as a source, and you are upgrading from an earlier version of DBMoto, you are required to update the stored procedures file on for DBMoto in your DB2 installation: ; • In the ServerFiles/DB2UDB folder under the DBMoto install folder, copy the new stored procedure library (name ends in 03) to the DB2 server. • In the DBMoto Management Center, edit the Connection Properties for the IBM DB2 connection, then edit the Transactional Log Type field to open the Setup Info dialog. • Click Verify and then Install to install the new stored procedures. • Fixed a problem that occurred during refresh from Oracle to Netezza when the source table had a Number field with a scale greater than zero and a record with a value that used all the allowed scale digits; • Fixed an InvalidCastException that occurred when refreshing an Oracle Number to a Netezza target; • Fixed a missing log file error that could occur when mirroring from Oracle after a log reset; • Modified the DBMoto logging to add a warning message when a source table is reorganized; • Fixed an issue that caused the MySQL log reader to go into an infinite loop; • Fixed a problem that sometimes gives rise to the exception "System.FormatException: Input string was not in a correct format." • Corrected the handling of missing update-after records. **Management Center** • Fixed graphical user interface problem in Source/Target Connection wizards for Informix connections, where Transaction Log window was not displayed to the correct scale. **API** • Fixed a null reference exception when removing a table from a connection.
7.1.6.2	**07/10/12**	• Added support for DBMoto Cloud Edition **Management Center/Data Replicator** • Added log-based replication support for DB2 V9.7 on HP-UX Itanium 11.31; • Fixed a problem where the Conflict Resolver Priority remained set when the priority changed to a value different from TargetServerWin. In this case the conflict resolver could fail with the following error: "A synchronization conflict could not be resolved because the script event Replication_onLateConflict has returned a null record." **API** • Fixed a problem in the API Library, where field mappings and schedules were not cleaned up when a replication was removed;

		• Improved trace information for the API calls; • Fixed a problem where Interface IInformixChangeDataCapture and classes LogMessage and HistoryMessage were not visible.
7.1.5.5	**05/24/12**	**Management Center** • Fixed a problem where log entries related to the group were not listed when when showing the history of a grouped replication. **Data Replicator** • Fixed an issue where the history showed an incorrect record count in synchronization; • Fixed a problem where Oracle 9i partitioned tables were not replicated in transactional mode; • Fixed a sporadic "Missing record AFTER for an UPDATE" error in trigger-based replication from IBM DB2; • Fixed the error "Input string was not in a correct format" when starting a synchronization; • Fixed error in synchronization in Microsoft SQL Server 2000 log-based replications: "System.Data.SqlClient.SqlException: Incorrect syntax near 'CAST'". • Fixed a problem with replications from an IBM DB2 for i (AS/400) source where mapping the RRN to a target primary key field fails on version V5R4 or higher. **API** • Fixed an index out of bounds exception that occurred when calling the StopReplicationManager API; • Fixed an issue with missing Verifier APIs in the 64-bit version.
7.1.4.7	**04/12/12**	**Management Center** • Introduced a "Destination ID" property for Oracle connections; • Changed the Mirroring Block Size default value to 0 (disabled) for SQLServer and Oracle, and 10000 for IBM DB2 for i (AS/400); • Fixed a problem where, using metadata on SQL Server, the following error could occur: "System.Data.SqlClient.SqlException: Could not find prepared statement with handle 4."; • Fixed an issue with SQL Server providing an incorrect last ID when read by date time; • Fixed an issue where the schema refresh function detected changes on table properties but did notcommit them to the metadata; • Fixed an error "Error reading from the log" which occurred when opening the Log Viewer. • Fixed an issue where the Metadata Explorer erroneously added a dummy node to the treeview if the attempt to create a default metadata connection failed. **Data Replicator** • Improved performance for DB2 LUW Log Reading replications;

		• Fixed a problem where the INFORMIX Log Reader Windows service stops when encountering rollback record in the log; • Fixed an issue where column headers were in the wrong order when using a target file connection; • Fixed a problem where an INFORMIX Log Reader replication would throw an exception when encountering datetime columns if the DBMONEY environment variable was set to comma; • Fixed an issue where the INFORMIX Log Reader Windows service crashed when there were many tables each with many columns; • Fixed some issues resolving synchronization conflicts; • Fixed an error that occurred when executing a refresh on IBM DB2 for i(AS/400) tables having column names starting with '@'; • Fixed a problem where INFORMIX mirroring replications did not replicate empty varchar and lvarchar correctly; • Fixed a problem for Oracle and SQL Server that occurred when mirroring in blocks failed to parse records with continuation flag; • Fixed some issues using synchronization with groups; • Fixed a sporadic NullReferenceException in the SQL Server log reader; • Fixed an issue for Sybase ASE trigger-based replications where transactions were processed in the wrong order; • Fixed an Oracle performance issue that occurred when opening the logminer with versions older than 10g; • Fixed a problem where mirroring with an IBM DB2 for i (AS/400) source always enabled an Arabic trace when 420 codepage was used; • Fixed an issue for IBM DB2 for i (AS/400) mirroring mode related to a System.Exception "Call to stored procedure CALL DBMOTOLIB.JRNSQNM VALUES FROM & TO not valid"; • Added an option for INFORMIX mirroring replications to "NOT turn off data capture/logging"; • Fixed an issue with INFORMIX sources where a replication could not be deleted after a table ID change was refreshed; • Fixed an issue where LOB fields larger than 32K were truncated in refresh replication to Oracle targets; • Fixed an issue In Sybase ASE trigger-based replications where an UPDATE on multiple records could produce a wrong result in the target; • Fixed an issue in Sybase ASE trigger-based replications where the cleaning of the log tables failed when the TID reached 20 digits; • Fixed an issue with refreshing a group with a File target where a new line was missing in the output file before the truncate table record.

		API • Fixed an issue with the AddTable API where an exception was thrown if a value was passed for catalog or schema for a DBMS that does not have catalogs or schemas; • Fixed an issue where the APILibrary sometimes fired the OnReplicationUpdated event even if the replication was not changed; • Fixed an issue to make the APILibrary group property MirroringInterval persistent.
7.1.3.5	**01/30/12**	**Common** • Added support for IBM i (AS/400) journal sequences greater than 10 digits; • DBMoto now automatically detect when a journal is configured to reset the sequence on a receiver change; • Introduced a Mirroring Block Size connection property for Microsoft SQL Server and Oracle. **APILibrary** • Added the method ITable.GetPrimaryKeys to the DBMoto APIs. It retrieves the list of primary key fields for a table. **Management Center** • Fixed an issue where an "Invalid character in the given encoding" error could happen while restoring metadata when the field description had non-ASCII characters. **Data Replicator** • Fixed an issue with Float(p,s) data type in MySQL trigger based replications; • Fixed a problem where an overflow exception was thrown refreshing an Oracle table when a NUMERIC field had a value exceeding 28 digits; • Fixed an Issue with the Informix Decimal date type that produced a sting truncation error; • Fixed an error that occurred when working with an AS400 V7R1 or higher with a table having a field named PRIOR. PRIOR has become a reserved word on these systems; • Fixed a problem with MySQL bulk-insert where regional settings use commas (',') as decimal separator; • Fixed an issue where Ritmo/i threw an exception "Index was outside the bounds of the array" while inserting on CHAR fields marked with CCSID 420 (arabic); • Fixed an issue where INFORMIX Log Reader replications do not replicate decimal values correctly; • Fixed an issue where, on a Netezza target, character fields containing the string value 'Null' were replicated as NULL. **Log Reader** • Fixed a problem where a SQL Server publication (and all the articles defined in it) was removed when deleting the last DBMoto replication that was using the publication;

		• Introduced a Microsoft SQL Server connection property to specify options for log-reading queries; • Fixed a problem where INFORMIX Log Reader replications did not replicate updated NULL values; • Fixed a problem where the INFORMIX Log Reader Windows Service failed on tables with more than 197 columns.
7.1.2.5	**10/31/11**	**Common** • Added "Encrypted User Password And Data" security mechanism for IBM DB2 UDB; • Added support for Silverlining as target connection. **APILibrary** • Implemented an APILibrary.IMetadata connection object and primary and secondary backup. Allowed adding and removing an IServer and an IMetadata object through the API; • Implemented an API to execute the replication verification; • Fixed a problem where it was not possible to call Enable, Disable or Clear if the replication had not been saved into the metadata; • Fixed a problem which occurred when creating a replication through API. The Start Time was not set; • Fixed an index out of range problem which occurred when using the API to call a RefreshCatalog on a table containing foreign keys; • Added trim for catalog and schema names retrieved when calling a RefreshCatalog using the API; • Removed a null reference exception which occurred when using the DefaultMetadata object if the server was not connected and the metadata list was empty; • Fixed the Resynch function by making it synchronous and synchronized multiple accesses to IReplication code from separate threads. **Management Center** • Added a feature to select tables to add to the metadata from an external CSV file; • In the DBMoto Verifier, improved comparison classes for date time data types; • Increased the maximum number of messages in the DBMoto log when saved to a database. The new maximum value is 1,000,000; • Fixed a problem with creating new metadata. The license for the current metadata object was null; • Fixed the custom restore replication feature. When the "Keep the original Transaction Read Point" check box was selected, the new replication did not save the Last Refresh Time and it would start with a refresh;

- Fixed a problem which occurred after creating a target table. The table could not be retrieved in the database because it was created with a different case, requiring the intervention of the user. The new table was added to the metadata but not to the Metadata Explorer treeview;
- Fixed a problem which occurred when removing a table (and subsequent fields and foreign keys.) The remove process could generate a 'collection modified' error;
- Fixed a problem which occurred when opening a SQLCommand panel with a user having BrowseConnection = false. If the login was not validated, an error was generated ('Cannot access a disposed object');
- Fixed a "collection was modified" error which occurred when removing tables and/or replications while exiting out of the Multiple Replications wizard;
- Fixed a null reference exception which occurred when creating a target table using drag & drop, and the drop node was a schema or catalog node not expanded but already loaded with sub-nodes;
- Fixed a problem with an error "There is already an open Data Reader associated with this command," generated when creating mirroring replications on SQL Server 2005.

Data Replicator

- Fixed a problem with bulk insert in MySQL. When primary key violations were generated, the Management Center would show no errors and an incorrect number of processed records.

Log Reader

- Added an ORDER clause option when creating SQLAnywhere triggers to manage multiple triggers per table;
- Added support for IBM DB2 UDB data types GRAPHIC(NChar) and VARGRAPHIC(NVarchar);
- IBM DB2 UDB Log Reader can now read changes from partitioned tables;
- Fixed an issue in mirroring from an IBM DB2 for i source where the last processed transaction was re-processed again after stopping and restarting the Data Replicator;
- Fixed a problem replicating from an Oracle RAC;
- Updated the program installed on IBM DB2 for i. The search for the last sequence is now performed on the current receiver rather than the current chain;
- Fixed a problem in the IBM DB2 UDB log reader where the replication removal did not turn off the data capture on the table;
- Fixed a problem which occurred when replicating to a File connection. The target table truncate was not recorded in the output file;
- Fixed a problem where the IBM DB2 Log Reader threw an exception when the tableid or tablespaceid number was negative.

7.1.1.1	**07/11/11**	**Common** • Introduced the "Max Commit Interval" property for Netezza target connections **APILibrary** • Added GetReceiversInUse API and SampleDeleteReceivers API sample.; • Fixed a Null Reference exception in the Data Replicator which occurred after stopping the trace from the Service Monitor. **Management Center** • Fixed a license check error generated when the license file was missing or contained an invalid license. (Error message: Collection cannot be null.); • Fixed a problem with custom restore replication. When the "Keep the original Transaction Read Point" checkbox was selected, the new replication didn't save the Last Refresh Time and it would start with a refresh; • Fixed a problem in the Replication Wizard and Multiple Replication Wizard where the list of groups in the Replication Type step was not ordered alphabetically; • Fixed a null reference exception that occurred when switching a replication using and IBM DB2 UDB source from refresh to mirroring; • Fixed an error that occurred with SQL Server replications when reading the last ID. After checking "Refresh Replication Objects" and "Maintain Current Sequence," new values such as the Article ID were not saved in the metadata. • Fixed a problem with the Verifier. It did not work properly when there were multiple primary keys in a different order compared to their ordinal column order. The ORDER BY statement did not contain all keys in the correct order. • Introduced LoadUniqueIndexes property to load unique indexes information when primary keys are not defined; • Fixed a problem for remote distributors. The connection CommandTimeout value was not passed to the distributor commands. • Fixed a problem that occurred when creating new MySQL, DB2UDB and File connections. The transactional setup was set to the default log setting instead of the Disabled setting. • Added a Reset button to the Log Setup dialog for Informix Log Reader to allow the user to reset the log position file to 0 so that the log reader service restarts reading from the latest/current log position. • Fixed an issue with the Oracle supplemental log setup. Changed the query to manage single quotes and case sensitivity. **Data Replicator** • Fixed a problem where mirroring to Netezza did not use transactions if the Isolation Level was set to "Unspecified" in the Replication Properties;

		• Fixed an error that occurred when writing to the MessageArgs field in the log table using ';' separator instead of '|'; • Fixed a NullReference error that occurred when the DB2 UDB Log Reader encountered tables with a negative table ID. **Informix Log Reader Window Service** • Changed the maximum number of records read to improve read performance; • Changed log position file by adding a new last valid LSN line to make recovery more accurate.
7.1.0.9	**05/27/11**	**Management Center:** • Fixed a problem in the DBMoto History Viewer that occurred while reading history when the log type was set to database; • Fixes a problem in the DBMoto Log Viewer,that occurred while reading the log when the log type was set to database.
7.1.0.8	**05/16/11**	**Common:** • Fixed a problem that occurred when creating a trigger-based replication on Informix. DBMoto gave an error if the source table had triggers already defined on it. Since Informix now supports multiple triggers per table, the error has been removed for Informix version 11.1 and up; • Added support for transaction log-based replication from Informix v.11.5 FC3 and later. **APILibrary:** • Implemented new set of APIs using a .NET framework object model instead of a procedural approach • Implemented APIs to read from the DBMoto log and history. **Management Center:** • Added a Replication Activity Viewer to monitor and analyze replication performance. Available for .NET Framework 3.5, SP1 only; • Modified most common message boxes to add the "Do not show this again" checkbox; • Fixed a problem with SQL Server nvarchar(max) data type. It was shown in the Object Browser as having precision 0 instead of -1. This produced an error in the creating target table wizard; • In the DBMoto Verifier, fixed a cast issue comparing GUID fields; • In the DBMoto Verifier, fixed a cast issue. When there is an invalid cast exception comparing two items, the value and data type is displayed; • In the DBMoto Verifier, provided a way to handle SQL Server cases when Time is returned as Varchar;

		• Fixed an issue in the Custom Restore Wizard in the group selection step. If a group was selected, an exception was thrown when clicking Next in the wizard; • Fixed an error in the Replication Wizard that would corrupt the metadata; • For Netezza targets, the staging table is now removed when dropping the replication; • Fixed an error in the Metadata Explorer that occurred when the Data Replicator settings were opened and then saved; • Fixed a problem in the Global Script when adding a new reference that contained space chars, each space would split the reference into more references, causing compile errors on the Global Script. Fixed by using double quotes.
7.0.4.10	**03/01/11**	**Common:** • Replications using Netezza.now handle money and smallmoney datatypes with a numeric datatype. **Data Replicator:** • Solved an issue with many-to-one replications when Netezza is the target.
7.0.4.9	**02/09/11**	**Common:** • Introduced queries for each database to get the current database timestamp; • Fixed a communication problem from client applications (ManagementCenter, APIs) to the Data Replicator. If the metadata connection was temporarily closed, the TCPIP connection was unavailable. **Management Center:** • Fixed an issue with the SQL Query tab. When opening a result set, the operation was very slow due to a refresh function that needed optimization; • Fixed a problem with catalog queries. In some cases DBMoto would not find a table in a database even when it was already created; • Fixed a problem with initial refresh in groups. When adding a replication to a group, if the only difference was the Initial Refresh status, the error message would say 'replication mode differs'; • Fixed an error in the Metadata Explorer. When listing all replications and selecting Enable/Disable of a single replication, if the replication was in a group, the other replications in the group would not change icon as a result of the status change; • Fixed an issue with the metadata. When the metadata database was stopped, the Management Center was not able to reestablish a connection to the TCP/IP server; • Fixed a null reference exception in the Replication Wizard (in the Logging Info step) when Oracle was used as source;

		• Fixed a problem in the Replication Properties dialog. When switching a replication having Oracle as a source from mirroring to refresh, the dialog would not change replication mode; • Fixed a problem that occurred when removing a server from the Metadata Explorer and adding a new server with the same name. A message would be displayed saying that a server with that name already existed. **Log Reader:** • Fixed an IndexOutOfRange exception mirroring from a MySQL source when an enum column is blank; • Fixed an ArgumentOutOfRangeException in mirroring mode from a MySQL source with transaction l og reading that occurred when the date in a datetime value had zero yearFixed an ArgumentOutOfRangeException in mirroring mode from a MySQL source with transaction log reading that occurred when the date in a datetime value had zero year.
7.0.4.8	**01/24/11**	**Common:** • For Netezza, added support for automatic detection of Primary Keys; • Added the "Case Sensitive Table Names" option in the MySQL Transactional Support configuration dialog; • For Oracle, added table-level supplemental log management; • Added the option "Fire Triggers" in the refresh options of a replication when the target connection is SQL Server. The option enables and disables triggers when the insert mode is set to bulk insert; • Introduced a version check when connecting from the Management Center or a client API application to the Server Agent. If the versions are not compatible, a message is shown and the connection is denied; • Fixed a problem where a license allowing synchronization between two database of the same type did not allow mirroring both ways. **API Library:** • Fixed a problem with API traces: for the way traces were initialized, they would not catch possible errors generated at the beginning of the Initialize API. **Management Center:** • Modified the Verifier so that each verification runs in a separate thread; • Modified the Verifier so that the "Save Results" feature now saves the data in the grids; • Modified the Verifier so that the 'Save Results…' button now saves the details, queries and the differences on the file specified; • Modified the Verifier to remove the "Open Results" feature;

		• Modified the Verifier to add "Verifier Sort Sequence Table" to the dynamic properties for connections to IBM DB2 for i (System/i, AS/400). The table should contain ordering rules for the data in the result set and is used by the Verifier to sort character values; • Modified the Verifier to set the Automatic Verification option to false by default in the config file; • Removed an incorrect warning message when running a 'Refresh Schema Information' operation on a table; • Modified the Multiple Replications Wizard "Set Replications" step where tables are selected in the grid. A "Source Type" column has been added to show the type of table (Table, View…); • Modified the Change Login operation. When it is called on a loaded server, all SQL Command panels and Log and History panels are forcefully closed. This is to avoid leaving opened panels on which the new login may not have permissions; • Fixed a problem where the DBMoto Verifier did not check the "Run Verification Automatically" setting; • Fixed a casting problem with TimeSpan data type during the verification; • Modified the SQL Command feature so that each SQL command runs in a separate thread; • Modified the LogViewer and HistoryViewer to use a dynamic paging mechanism. The new page is loaded when the user scrolls up and down the list; • Fixed a problem that occurred when altering a table associated to a trigger-based replication; • Fixed a problem that occurred when running the Verifier. Huge tables could cause the first compare to stop due to an error when the second one was started; • Fixed a problem with the SQL Command tab regarding error messages in a separate dialog. Error messages are now all reported in the Results pane; • In the SQL Command tab, disabled the Run button during execution to avoid multiple clicks on it; • Introduced new target connection properties: Default Insert Mode, Default Block Size and Default Isolation Level; every new replication will get the values of its related properties from the target connection default properties; • Improved DBMoto Verifier presentation of byte values in the data grid: System.Byte[] data is now written in hex format; • Fixed some minor problems on the Data Replicator Settings and the Options dialog. (Numeric fields set to blank would throw an exception); • Fixed a NullReferenceException when renaming a server node in the Metadata Explorer; • Fixed a problem in the Create Target Table wizard that occurred when creating a table on Oracle--a 'NUMBER' type with null precision was not accepted; • Fixed a problem in the Create Target Table wizard that occurred when creating a table on Netezza--a 'numeric' type with null precision was not accepted.

		Log Reader: • Implemented the Case Sensitive Table Names option. **Data Replicator:** • Fixed an error processing transactions from an IBM DB2 for i source which was reported in the log as "Writing to compression stream is not supported"; • Fixed a problem where, in a mirroring replication with an IBM DB2 for i source and a Netezza target, a Clear operation on the source table generated an error in the log. The message was a general "Error during mirroring" with no further explanation; • Fixed a problem where a replication could get disabled when a error occurred while connecting to the source database, (e.g.,a timeout); • Fixed an IndexOutOnRangeException which was thrown when comparing records having a binary field. The exception could happen in synchronization while detecting or resolving a conflict; • Added support for the IBM DB2 for i CLEAR operation during mirroring. The operation translates into a truncation of the target table; • Fixed the error "ORA-01727: numeric precision specifier is out of range (1 to 38)" when replicating to an Oracle NUMBER field; • Fixed a problem where an ArgumentException was thrown during a refresh to a Netezza target table when the target table name contained lower-case or special characters; • Fixed a problem where sometimes after starting and stopping replications with the Data Replicator running, a replication could remain blocked in a stopping state; • Fixed a problem that occurred when calling the Replication_onError event. The replication/group name parameter was set to null in case of grouped replication; • Fixed a problem that occurred with PostgreSQL as target. On a "time without timezone" column, a cast error was generated during replication; • Fixed a problem with Oracle that occurred when using bulk insert. If errors were generated, the count of correctly inserted records (both in the log and in the monitor) was incorrect.
7.0.3.3	**09/28/10**	**Common:** • Fixed an error introduced in version 7 on Oracle 8.1.7 where the catalog query to retrieve the column list would generate an error; • Fixed an error that occurred when calculating the current DBMoto log file in case of logging on multiple files: if more than one file was most recently modified at the same time (seconds are not saved in the LastWriteTime property of a file), the first in the list was considered the current log file.

		Management Center: • Fixed an error that occurred when opening the Log Viewer or the History Viewer. Any subsequent operation on the metadata would throw a closed connection exception; • Fixed a problem with the User Permissions dialog where some of the permissions were not saved when closing the dialog; • Fixed a problem in the Custom Restore wizard where transactional setup on the involved connections was not checked when creating replications, which could generate a null reference exception; • Fixed an exception in the Custom Restore Wizard that occurred when importing Unmapped Fields and Non Critical fields; • Fixed a problem that occurred when opening the SQL Query tab with a login. The dialog would prompt a user name and password. This was not correct when the dialog was opened with the permission 'BrowseConnection' set to false; • Fixed a null reference exception that occurred when opening the SQL Query tab and clicking Cancel in the login dialog; • Fixed a problem that occurred when editing the Data Replicator Options. The dialog would save all properties and additionally overwrite options that were manually set in the dbmoto.server.config file. **Log Reader:** • Fixed an IndexOutOfRangeException error in a group when one or more replications have non-critical fields set; • Fixed a problem where replications from IBM i (AS/400) source connections could stop replicating or lose data if a receiver was changed with the *RESET option; • Fixed an issue that occurred when reading the IBM i (AS/400) transaction log which wrongly reported the error message "Missing AFTER image for an UPDATE". It could be followed by a NullReferenceException.
7.0.2.9	**08/23/10**	**Common:** • Added a new 64bit build to handle ODBC direct calls in a Windows 64bit installation; • Fixed a problem in the IBM DB2 for i (AS/400) DBMoto library when replicating a group of 19 replications. The DBMOTOLIB version has been updated to 7.1.x; • Fixed a problem reading timestamp values from IBM DB2 for i when the CCSID was 65535. The error was "An exception was thrown during the data conversion function. Cannot find a converter for ccsid 65535". **Management Center:** • Fixed an issue that occurred when using the Refresh Filter editor. It would run a compile check even if it was editing SQL script and not DBMoto script ;

- Fixed an error that occurred when opening an Execute Command dialog on the metadata connection and running any query, then closing it. Any subsequent operation on that metadata would throw a connection error because the connection to the metadata was closed;
- Fixed an "Index out of bounds" exception that occurred when creating a new group using the wizard, then selecting Group Properties and adding a replication to the group;
- Fixed a null object reference exception in the Metadata Custom Restore wizard. It was caused by replication names containing blanks;
- Fixed an issue with the Trace Replications dialog used to filter replications for the Data Replicator trace. The value of this property was saved but not correctly displayed in the dialog;
- Fixed an issue that occurred when reading the Transaction ID from the Replication Properties dialog. The CommitTID value needed to be cleared out to stop DBMoto from picking up transactions from a previously saved transaction read point;
- Fixed some help context dialogs. Some dialogs would not open the help manual using the F1 button;
- Fixed an index out of bounds exception that occurred when refreshing schema information on a table which was modified using an ALTER TABLE command;
- Fixed an issue in the Metadata Custom Restore Wizard and the Export Replications Wizard where the information about Unmapped fields and Non Critical fields was not restored in the newly imported replications;
- Introduced automatic deletion of a staging table, when the refresh schema information is invoked, for a Netezza table defined in a transactional replication where the table structure has been modified;
- Introduced a new option in the Data Replicator Settings dialog to define the DBMoto temporary files location (by default created in the SyncCache folder);
- Fixed an issue that occurred when changing path to the DBMoto log or trace settings (only possible by manually editing the dbmoto.config and dbmoto.server.config files.) The path was overwritten with the default location when the options dialog was opened and saved;
- Fixed a null reference exception generated in the Custom Restore Wizard when a replication name to be imported contained blanks at the end.

Log Reader:

- Fixed an issue with a missing transaction in the IBM DB2 for i log reader.

Data Replicator:

- Fixed a problem where a Null Reference could occur when starting the Data Replicator;
- Fixed an index out of bounds exception in synchronization mode;
- Fixed an InvalidCastException in synchronization mode.

7.0.0.18	**06/21/10**	**Common:** • Fixed a problem in the metadata Custom Restore Wizard. Replications associated with a group were not loaded with the group information if the replication status was disabled; • Fixed a problem in the metadata Custom Restore Wizard. When restoring a replication with IBM DB2 for i as a source, if the checkbox to keep the current transaction ID was selected, the receiver library and name associated with the replication was not imported; • Added support for synchronization between more than two tables; • Introduced conflict priority for groups with multi-table synchronizations; • Added features to support secure access to DBMoto and user management; • Introduced live backup connections for metadata, to allow saving metadata on multiple databases simultaneously; • Added support to start and stop tracing without having to restart the Data Replicator; • Added a command timeout property to the metadata connection; • Added support for Microsoft Access 2007; • Added support for Netezza as target database; • Added support for DB2 LUW transactional replications using transaction logs; • Introduced support for the new .NET provider for Unify SQLBase (Unify.SQLBase.Data.dll); • Added support for GRAPHIC data types on DB2 LUW; • Added support for Oracle data types binary_float and binary_double; • Added support for the SQLServer 2008 data type DatetimeOffset; • Added support for the XML data type for SQL Server 2008 and Oracle version 10 or later; • Added a custom mapping script function with isNullable and isPrimaryKey parameters; • Added an option in the XML configuration file for databases that do not allow nullable fields as primary keys, to create target tables according to this setting; • Upgraded the metadata to encrypt the Connection String related to the log on database option. Previously it was saved unencrypted; • Changed signature to the events onBeforeMirroring and onAfterMirroring with a new Boolean parameter that indicates the direction of the mirroring in case they are called in synchronization; • Added a hyperlink on known errors to KB articles for common problems in the Management Center and Log Viewer; • Added support for using filters for replications and groups in the trace options; • Changed DBMoto license to include also the maximum number of clients connected; • Introduced a new AddLog API that uses an enumerated parameter to define Information, Warning or Error (the previous definition is still valid but obsolete). Added a further AddLog API that is able to printout the list of field values of an IRecord parameter; • Inserted an “Imports System” statement by default in all global and replication scripts to avoid common compile errors;

		• Removed a limitation that forced the replications in a group to have same values for the following properties: Refresh recovery, Insert Mode, Block size and Isolation Level; • The bulk insert for the Microsoft SQL Server OLEDB provider, also known as Fast Load, is no longer supported; • Fixed a problem where the length of nchar and nvarchar fields for SQL Server and Oracle was shown doubled in the EM Object browser; • Fixed catalog issues while loading metadata tables (the casing was ignored and in some DBMSs such as MySQL it was causing the application to ignore the existence of metadata tables); • Fixed an error on Foreign Keys tables when loading old metadata sets; • Fixed a problem in the metadata to avoid overflow errors when saving the setting MailRcptTo with a list of addresses with size greater than 50; • Fixed a problem in themetadata to avoid overflow errors when saving the Connection String and Connection Params setting related to the log on database option, when their size was greater than 50; • Fixed an issue with Oracle 64-bit where in some cases the length of character fields was doubled; • Fixed an issue with the GetReceiversInUse script API, receivers in use were returned even if defined in Refresh replications; • Fixed an issue where Oracle floating point data type NUMBER was in some cases wrongly mapped to a NUMERIC(38, 0) when creating target tables. **Management Center:** • Added the DBMoto Verifier tool to compare source and target tables after replication; • Added a filter in the Select Tables dialog to open only selected catalogs, schemas and tables depending on the filter defined; • Introduced multi-selection capability in Fields Mapping screens to edit more than one mapping at the same time; • Added support for creating catalog and schema levels within the Custom Restore wizard, instead of requiring that they are already defined in the metadata; • Added support for importing group definitions in Custom Restore wizard; • Allowed the Custom Restore wizard to load replications to the target server without changing the current id (useful when backing up and testing a situation in a different environment); • Enabled multiple editing to save the same script to multiple replications; • Added new properties to the script object: connection name, table name, replication name, group name; • Added a main window progress bar useful when running slow operations. When selecting tables from a database, the progress bar is made visible and an Abort button is shown; • Added support for changing the table CCSID (IBM DB2) on multiple selections;

		• Added a Save Results menu item in the SQL Query tab; • Improved the readability of group properties in the Property pane by displaying a list of replication properties with aggregated values when a group is selected; • Introduced an 'Add to Group' and 'Remove from Group' menu item in the context-sensitive menu for replications; • Introduced an 'Enable Group' menu item in the context-sensitive menu for replications to allow editing of multiple replications with a single click; • Added a Locate item to the Replication Browser toolbar to allow, from a selected replication, access to its source or target connection, catalog, schema and table; • Added a checkbox to the Transaction Read Point dialog (accessible via Replication Properties), to refresh the transactional replication objects without reading the new ID. This is useful when transactional objects have been removed and need to be created again, but the replication has to start from the same transaction ID that was set in the metadata; • Introduced a check to verify if the group insert mode and isolation level are compatible with the target connection of the replication when assigning a replication to a group. If not, the replication will keep its current insert mode and isolation level; • Moved Data Replicator Options from the Metadata menu to the Server menu, as they pertain logically to the server; • Modified the Log Viewer to show the current log file in italics to distinguish it from other log files; • Fixed an issue with running the "Create Multiple Replications Wizard" from Oracle to MySQL. On MySQL if the target tables already existed, the tables on the target database were not found when searched using upper case names; • Fixed a problem that occurred when switching from refresh mode to synchronization mode in the Replication Properties dialog. The transactional objects needed for the replication were only created in the source to target direction but not in the target to source direction; • Fixed a problem that occurred when changing the IP address of a source or target connection. If the connection had transactional support enabled, it was possible to leave the metadata in an incorrect state, having connection properties pointing to one server and transactional objects pointing to another server. Added a warning when the data source property is changed and transactional support is enabled; • Fixed a visualization issue in the Property Window. When a trigger-based replication for SQL Server was selected, the properties list showed publisher and article information instead; • Fixed an issue that occurred while running a catalog refresh on a catalog that has tables currently replicated. Previously, you would get a warning with an OK button for each one of the tables; • Fixed an issue with the history status. When a group was refreshing, the History Status became "success" while still refreshing, instead of waiting to the end of the operation;

		• Fixed an issue with the history status. A replication in a group that had been stopped dipslayed the history flag as Stopped, but when the group was refreshed again, the replication history status stayed Stopped; • Fixed the order of items in lists in the Object Browser, Replication Browser and Replication Monitor. When a sorting order is selected, the order is maintained when switching to a different node; • Fixed an issue in the Replication Monitor. Before starting the refresh, the replications were displayed as 100% completed; • Fixed a problem that occurred when restoring metadata on a SQLBase database; • Fixed a replication visualization issue. When two replications with different schedules were put into two different groups and then edited together, the schedule would disappear; • Fixed an issue in the SQL Query tab. The resultset and command are now closed when the query is over to avoid command timeout errors if the dialog is left open for too long; • Fixed an issue with the Unmapped columns and the Non-critical columns on the mappings definition of a replication. If the source or target table's structure was altered, an exception was generated when loading the mapping dialog; • Fixed a problem that occurred when making changes to a replication within a group (for example, the description.) After pressing OK, you would get thrown out of the group into the replications screen. Now the replication and the group you just configured is displayed; • Fixed a problem where the DBMoto Management Center would open in maximized mode and didn't save the form size from when it was last opened; • Fixed a problem where, in the DBMoto Management Center, the content of the Object Browser, Replication Browser and Replication Monitor did not refresh when a metadata was restored; • Fixed a problem in the Metadata Explorer. If a replication was put into a group or removed from a group, its icon in the treeview did not change; • Fixed a problem in Metadata Explorer. When clicking on empty nodes (i.e. "Sources, "Targets") cross metadatas, the replication content did not get refreshed; • Fixed an issue that occurred while creating a trigger-based replication on a database that allows only one trigger of the same type on the same table (e.g. Informix.) Now we verify if there is already a trigger defined on the source table and, if there is, report an error; • Fixed an issue where creating a mirroring or synchronization replication on an Informix source table containing a field with data type decimal(n), decimal, numeric(n) or numeric resulted in an error. **Log Reader:** • Fixed an error in trigger-based replications where the following message occurred: "A record is missing in the log table";

		• Introduced commitment control for IBM DB2 for i as a source. Only committed transactions are replicated to the target; • Added support for Journals set to minimized entry data as field boundary (MINENTDTA(*FLDBDY)); • Added a warning message in the DBMoto log when an RGZPFM (reorganized physical member) is executed on a System/i source table; • Included support for the TIMESTAMP alias name of the TIMESTMP data type for IBM DB2 for i; • Added support for save points in IBM DB2 for i and Oracle; • Fixed a problem that occurred when mirroring a SQL Server. A Float data type could result in a loss of decimal digits precision; • Fixed an error that occurred when creating a mirroring/synchronization replication on a table containing auto-incremental fields when a trigger based replication is set; • Fixed an error that occurred when creating a mirroring/synch replication on an Informix table with one or more datetime fields; • Fixed a problem where IBM DB2 for i date values with a 2 digits year format were sometimes assigned to the wrong century; • Fixed an issue with Oracle where we performed restricted verification when checking for SUPPLEMENTAL LOGs enabled. Now we also check if the primary key clause has been included in the supplemental log; • Fixed an issue with a NullReferenceException that occurred when reading a record from the Oracle Log Reader while an Oracle INTERNAL operation was processed; • Fixed error ORA-310 that could happen while reading the Oracle transaction log; • Fixed an issue with a missing transaction timestamp in mirroring and synchronization replication from SQL Server. **Data Replicator:** • Introduced support for the BIT data type on default size (1) in MySQL bulk insert; • Introduced 'Insert On Update' feature to define whether on an update failure DBMoto will run an insert statement; • Added a message in the DBMoto log when a mirroring/synch resumes after a scheduled refresh; • Added a Refresh Interval parameter to configure how long to wait before retrying a refresh. If a refresh replication could not start for some critical error (not disabling the replication), refresh was re-tried on each subsequent mirroring interval; • Modified the error message that is generated when there is an exception while evaluating a mapping expression. The target column name mapped to this expression is now displayed; • Modified the log message for errors processing INSERT transactions in mirroring. It now shows the transaction ID and transaction timestamp of the failing transaction;

		• Fixed a problem where, in some cases, the log message for errors processing transactions in synchronization did not show the replication name; • Fixed an error in Replication History. It used to update the history even when last date of refresh/mirroring was not changed; • Fixed an issue where a scheduled recurrent refresh of a group did not truncate the target table; • Fixed an issue happening when altering (by adding or removing columns) a table used in replication with unmapped fields set; • Fixed an issue in Ritmo/i that occurred when refreshing a table containing inline BLOBs; • Fixed a problem where, for a SQL Server target with bulk insert set, NULL values were converted to default values when running refresh; • Fixed a problem where, for replications configured in bulk insert mode or simulated bulk mode, the Refresh_onAfterRefresh event could be called before the last bulk block was processed; • Fixed an issue where metadata tables could become locked following a recovery from an error accessing the metadata. This lock prevented the Mangement Center from opening the metadata, resulting in a timeout error; • Fixed a problem where, in synchronization mode, the replication properties InternaSourceConnection and InternalTargetConnection were not returning the correct connection; • Fixed an issue where precision was lost with SQLServer float data type; • The data type float was internally mapped to a Float instead of a Double. Fixed the SQLServer.xml file; • Fixed a problem replicating from/to a Sybase numeric and decimal with precision greater than 28; • Fixed a problem where an error was reported reading the Oracle transaction log when the record contained a LONG data type filed; • LONG was mapped to an internal logvarbinary instead of longvarchar; • Fixed an issue with the temporary files in the SyncCache folder while running synchronization. The file could not be accessed because was opened by another process. **Service Monitor:** • Added monitoring for the DBMoto Server Agent, with menu items for starting and stopping the agent; • Added a menu item to launch the Management Center; • Fixed a problem where, when connected remotely and the Service Monitor was running, on log off the Service Monitor was shut down and it would show the warning 'The data replicator will keep running even if the Service Monitor is stopped'.