



Informatica Mass Ingestion January 2023 Release Notes

© Copyright Informatica LLC 2019, 2023

Contents

Common.	2
Fixed issues.	2
Known issues.	2
Mass Ingestion Applications.	3
Fixed issues.	3
Known issues.	3
Mass Ingestion Databases.	4
Fixed issues.	4
Known issues.	9
Informatica Global Customer Support.	14

Read the *Mass Ingestion Release Notes* to learn about fixed issues and known issues in the January 2023 release of the Informatica Intelligent Cloud ServicesSM Mass Ingestion service, which includes Mass Ingestion Applications, Mass Ingestion Databases, Mass Ingestion Files, and Mass Ingestion Streaming.

For information about new features and enhancements, see *Mass Ingestion What's New*.

Common

Fixed issues

The following table describes fixed issues that apply to two or more ingestion types. Not all monthly releases include fixed issues. Note that the (*month, year*) value in each issue description represents the release in which the issue was fixed.

CR	Description
DBMI-11705	While you create an application ingestion or database ingestion task in the task wizard, the runtime environment options are not listed if read permissions are not assigned to the Secure Agent Group. (November 2022 release)
DBMI-11347	<p>If you modify a source column after you run an application ingestion or database ingestion incremental load or combined initial and incremental load job that has a Google Big Query target, the job ends abnormally with the following error:</p> <pre>Could not execute the following SQL query : 'ALTER TABLE `schema`.`table` MODIFY `column` STRING;...because of following error: [Simba][BigQueryJDBCdriver](100032) Error executing query job. Message: 400 Bad Request</pre> <p>(October 2022 release)</p>

Known issues

The following table describes known issues that apply to two or more ingestion types. Note that the (*month, year*) value in each issue description represents the release in which the issue was found.

CR	Description
CCON-47840	<p>If you try to test a connection to a Databricks Delta source or target using a Secure Agent on Windows, the test fails with an error message that states Windows is not a supported operating system. However, the connection works when you create the ingestion task and run the associated job provided that the connection properties do not include a SQL Endpoint JDBC URL value. This issue applies to all ingestion types.</p> <p>Workaround: To test the connection, use a Secure Agent that runs on Linux. Also ensure the SQL Endpoint JDBC URL value is specified in the connection properties. After the test is completed, remove the SQL Endpoint JDBC URL value to be able to use the connection when creating a task or running the job. (November 2022 release)</p>
DBMI-12113	<p>Application ingestion jobs and database ingestion jobs that process a large number of source columns and many rows of data for an initial load might fail with an error message that reports the GC overhead limit was exceeded. For example, this problem can occur with a SQL Server source column size of 31 bytes and 100000 rows of data to be unloaded to a SQL Server or Snowflake target.</p> <p>Workaround: To resolve this issue, increase the Secure Agent's heap size to a value such as 32 GB. (January 2023 release)</p>

CR	Description
DBMI-9680	If application ingestion or database ingestion jobs capture empty fields or columns from sources, the corresponding fields on the target either will be empty for initial load jobs or contain nulls and empty strings for incremental load jobs. Workaround: None. (November 2021 release)
DBMI-7107	When you deploy an application ingestion or database ingestion task, the target objects are generated with a structure that matches the current schema of the source objects. If schema drift causes the structure of the source objects to change before you run the job for the first time, the source structure changes are not detected. The changes are detected only after the job runs for the first time. Schema drift changes include adding, removing, altering, or renaming columns. Workaround: Do not change the source schema until after the first run of the job. (October 2021 release)

Mass Ingestion Applications

Fixed issues

The following table describes recent Mass Ingestion Applications fixed issues. Not all monthly releases include fixed issues. Note that the (*month, year*) value in each issue description represents the release in which the issue was fixed.

CR	Description
DBMI-7179	When you deploy an application ingestion job that is configured to load a large number of SAP ECC objects to an Amazon Redshift target, the deployment might fail. (November 2022 release)

Known issues

The following table describes Mass Ingestion Applications known issues. Note that the (*month, year*) value in each issue description represents the release in which the issue was found.

CR	Description
AIN-3849	When you create application ingestion initial load tasks that have a Workday source, if you select the Workday Report-as-a-Service (RaaS) API to extract source data from custom objects and fields through a single report and specify a URL in the Report Name or URL field, the associated ingestion jobs will fail with the following error: Wildcard expression contains unescaped .(dot). Unescaped dot should be used only as delimiter between object names Workaround: Specify a report name instead of a URL. (January 2023 release)
AIN-3695	Application ingestion initial load jobs that have a Workday source and use the Workday Report-as-a-Service (RaaS) API to extract source data from custom fields through a custom report fail if you previously created the custom report in Workday with prompt default fields marked as required. Workaround: None. (January 2023 release)

CR	Description
AIN-3437	In incremental load jobs with cloud data warehouse targets, when a schema change occurs on a source object, the subtask associated with the object fails even though the job is configured to replicate source schema changes on the target. This issue occurs in the jobs that are configured for NetSuite sources. (September 2022 release)
DBMI 11800	Application ingestion jobs configured to ingest data from Adobe Analytics or Google Analytics sources to Oracle targets might fail in the following scenarios: <ul style="list-style-type: none"> - When a source object contains a character type but does not contain a primary key constraint. - When a source character type is part of a primary key constraint. (November 2022 release)
DBMI-10668	Application ingestion jobs configured for Amazon Redshift targets fail if a source object contains a value of string datatype with a precision greater than 65000. (June 2022 release)
DBMI-7323	Application ingestion jobs configured for SAP ECC sources do not propagate objects that contain blank values if the column containing the blank value is mapped to a primary key on the target. (October 2021 release)
DBMI-7067	When an application ingestion job writes a value of the STRING data type to a Google BigQuery table, irrespective of the length of the source value, the target value is set to the maximum supported length. (October 2021 release)

Mass Ingestion Databases

Fixed issues

The following table describes recent Mass Ingestion Databases fixed issues. Not all monthly releases include fixed issues. Note that the (*month year*) value in each issue description represents the release in which the issue was fixed.

CR	Description
DBMI-12156	Database ingestion incremental load jobs that have a Snowflake target and use Audit apply mode with the Add Operation Sequence option might write duplicate sequence values to the audit table on the target, instead of writing a unique, ascending sequence value for every change operation. (January 2023 release)
DBMI-12070	When you try to run a database ingestion job that uses the mock Sample Oracle Connection or Sample SQL Server Connection as the source connection, the job might fail with the following error: <pre>No source connection config helper defined for type</pre> (January 2023 release)
DBMI-11944	In a database ingestion task with an RDS for Oracle source, if you use the custom property <code>pxw.cdcreader.oracle.option.logArchiveWait</code> to specify the number of seconds that Mass Ingestion Databases waits for the copy of the log to become available as a new archived redo log, your specified value changes to 300. (January 2023 release)

CR	Description
DBMI-11869	<p>Database ingestion jobs that have a Microsoft Azure Synapse Analytics target might fail intermittently while establishing a secure connection. The following error message is issued:</p> <pre>Failed to execute target queries. Error: The driver could not establish a secure connection to SQL Server by using Secure Sockets Layer (SSL) encryption. Error: "Unexpected rethrowing".</pre> <p>(January 2023 release)</p>
DBMI-11792	<p>Database ingestion jobs that have a Db2 for z/OS sources might incorrectly calculate the fixed row length. The log data is truncated when most columns are null, and an update occurs that causes the log row to expand beyond the computed length. (November 2022 release)</p>
DBMI-11771	<p>Database ingestion jobs that have an Oracle source and a function-based index as part of the source table might fail while validating the index metadata. Mass Ingestion Databases validates the index name against a source table column name. The job fails when a matching column name does not exist. (January 2023 release)</p>
DBMI-11761	<p>When a database ingestion incremental load job that has a Db2 for z/OS source is running, the after image might be truncated if the data changes are in a column that does not use reordered row format. (November 2022 release)</p>
DBMI-11754	<p>Database ingestion incremental load or combined initial and incremental load jobs that have a SQL Server source might fail while change data reading from SQL Server CDC tables is initializing. The following message is issued:</p> <pre>221020 130308 WIN64 36912 PWX-36749 Program internal error. Prog="SSLRCap::ReadStartupCDC". Line=759. Desc="Failed to initialize CDC reader."</pre> <p>(November 2022 release)</p>
DBMI-11728	<p>Database ingestion incremental load or combined initial and incremental load jobs that have a SQL Server or PostgreSQL source and process a large amount of transactional DML operations might fail with the message:</p> <pre>ThreadMain: Wait event has signalled without clearance.</pre> <p>This error can occur when event messages are forced to queue until Mass Ingestion Databases can accept new events because of a problem with the wait mechanism. (November 2022 release)</p>
DBMI-11709	<p>Database ingestion incremental load or combined initial and incremental load jobs that have a SQL Server source might produce a mismatch between the source and target data because DML records are not being replicated to the target. This error occurs when a recovery scenario causes truncated log reads. (November 2022 release)</p>
DBMI-11708	<p>While you run multiple database ingestion jobs, Mass Ingestion Databases might fail to register heart-beat responses from individual ServiceGroup members. As a result, frequent restarts of Database Ingestion agent services might occur. (November 2022 release)</p>
DBMI-11699	<p>The log collection tool fails when looking for the Database Ingestion agent's dbmiagent.properties properties file because the path name contains an incorrect separator. For example:</p> <pre>/app01/infaagent/apps/Database_Ingestion/conf/active/services/dbmi_agent \dbmiagent.properties</pre> <p>(January 2023 release)</p>

CR	Description
DBMI-11644	Database ingestion incremental load jobs that have been stopped continue to have the status of Stopping. The job status does not change to Stopped and the task log shows no recognition of the stop request. (November 2022 release)
DBMI-11633	A database ingestion job that has a SQL Server source does not replicate data to the target successfully if the source table contains 14 or more primary key columns. (November 2022 release)
DBMI-11573	After a database ingestion job that has a Db2 for i source fails, the job retains the Running status. (November 2022 release)
DBMI-11564	For database ingestion initial load tasks that have a Databricks Delta or Snowflake target, the Apply Mode field is available on the Target page in the task wizard even though the field applies only to incremental load and combined initial and incremental load tasks. (November 2022 release)
DBMI-11560	If the Database Ingestion agent service marks a database ingestion job with the status of AWAITING_AUTO_RESTART for the reason CONTAINER_FAILED while also sending a request for the next task, the agent service encounters a status update conflict when it gets the same job again for status updating. As a result, the job status changes to On Hold. If you try to stop or abort the job, it hangs in the Stopping or Aborting status. (October 2022 release)
DBMI-11559	Database ingestion incremental load jobs that have a file-based target, such as Amazon S3, Google Cloud Storage, or Microsoft Azure Data Lake Storage, might fail with the following exception: Unexpected exception : Could not locate persistent table state entry for table [table_name] (January 2023 release)
DBMI-11534, DBMI-11536	Database ingestion incremental load jobs that have a MySQL 5.7 source might fail with the following error: DBCapi Connect(), SQL Exception, 0 08S01 Communications link failure The last packet sent successfully to the server was 0 milliseconds ago. The driver has not received any packets from the server.. Caused by: JDBCapi Connect(), SQL Exception, 0 08S01 Communications link failure (January 2023 release)
DBMI-11457	If you stop a database ingestion job and then abort it, the job retains the status of Stopping or Aborting until the Secure Agent is restarted. (October 2022 release)
DBMI-11454	If a database ingestion incremental load or combined initial and incremental load job does not perform an Update operation when all columns have the same value, Mass Ingestion Databases reports excessive "Skipping update operation" warning messages in the log. (October 2022 release)
DBMI-11424	When you run a database ingestion initial load task that has a MongoDB source with a collection of more than 100,000 documents and an Amazon S3 target, not all of the documents in the collection are unloaded to the target. (December 2022 release)
DBMI-11373	When defining a database ingestion initial load task with a SQL Server source, if you select the Include LOBs option for the source and then select an Oracle target, the Include LOBs option is deselected and becomes unavailable. (October 2022 release)

CR	Description
DBMI-11370	<p>When a source table is not valid for change data capture, database ingestion incremental load and combined initial and incremental load jobs end with an error message about metadata columns instead of issuing the following warning message and then continuing:</p> <pre>PWX-36740 CdcReader <sslr> Capture processing was stopped on table dbo.D_Lead_Trace.</pre> <p>(October 2022 release)</p>
DBMI-11308	<p>When database ingestion jobs with a Kafka target run, the performance of writing data to the target might be degraded if Kafka producer threads are blocked on performing a callback. Also, the wait time counter in the following message is not reset:</p> <pre>[CDCPUB_13015] The process [Callback Helper] with the thread ID [identifier] is waiting for incoming data. Current wait duration hh:mm:ss.</pre> <p>(October 2022 release)</p>
DBMI-11237	<p>When deploying a database ingestion task for source tables that have backward slash (\) or hyphen (-) characters in a table name with defined target renaming rules, Google Big Query does not create the table and the task does not deploy. (January 2023 release)</p>
DBMI-11193	<p>A database ingestion incremental load or combined initial and incremental load job that has an Oracle source and an Oracle target might fail with the following error when an Update operation occurs on the source:</p> <pre>Error: [informatica][Oracle JDBC Driver][Oracle]ORA-00932: inconsistent datatypes: expected DATE got NUMBER</pre> <p>The Update is not written to the target. (October 2022 release)</p>
DBMI-11182	<p>If a schedule is assigned to a database ingestion task and the task user has been assigned specific permissions, or if the task is located in a project or project folder to which permissions have been assigned, the task does not start at the scheduled time. The following error is reported in the log:</p> <pre>{"errorCode":403,"errorMessage":"User autotest does not have permissions to perform action.", "additionalErrorMessages":null}</pre> <p>(October 2022 release)</p>
DBMI-11071	<p>Database ingestion jobs with SAP HANA sources that include decimal columns might use an incorrect precision and scale for these columns in the schema. The jobs always use a precision of 38 for decimal columns and write the actual column precision values as scale values. (October 2022 release)</p>
DBMI-11013	<p>If you deploy a task with a source table that has indexes built on unsupported data types, Mass Ingestion Databases does not filter out these indexes. Mass Ingestion Databases reports a schema validation error and deployment fails. (November 2022 release)</p>
DBMI-10954	<p>When you create a database ingestion task that has an Oracle source with NUMBER columns defined with a negative scale value, the preview of columns on the Source page of the task wizard omits the negative scale. For example, the preview might show NUMBER(30) instead of NUMBER(30,-12). If you use the incorrect data type specification from the preview in a custom mapping, the intended custom target data type is not reflected on the target. (October 2022 release)</p>

CR	Description
DBMI-10799	If a database ingestion initial and incremental load job stops because the Schema Drift Add Column or Modify Column option is set to Stop Job and then you resume the job with the Resume with Options > Resync option, the job resumes and the target tables are re-created with the DDL changes but none of the unloaded data appears on the target. (October 2022 release)
DBMI-10792	Database ingestion combined initial and incremental load jobs with a Snowflake target might encounter intermittent PUT exceptions during the unload phase. The retry mechanism is not triggered during the unload phase. (October 2022 release)
DBMI-10790	<p>If a database ingestion combined initial and incremental load job with a SQL Server source and Oracle target stops because the schema drift Drop Column option is set to Stop Job and then you resume the job with Resume With Options > Resync option, the target is not re-created and data is not loaded during the initial load phase. The job fails with the following error:</p> <pre>LEVEL_1 trace - Process <process> with thread ID <identifier> failed with unexpected exception: Index: 14, Size: 14</pre> <p>(October 2022 release)</p>
DBMI-10649	If the schema drift Rename Column option is set to Replicate for a database ingestion incremental load or combined load task that has an Oracle source and an Oracle or Snowflake target and then a Rename Column operation occurs on the source, the column with the old name is dropped and a column with the new name is added. As a result, the data that was captured under the old column name is lost. (October 2022 release)
DBMI-10323	Database ingestion incremental load jobs that have SAP HANA sources might fail if they're replicating data to VARBINARY columns on Amazon Redshift, Databricks Delta, Google BigQuery, or Microsoft Azure Synapse Analytics targets. (October 2022 release)
DBMI-9881	If a database ingestion incremental load job with a SQL Server source and Oracle target stops because the schema drift Drop Column option is set to Stop Job or Stop Table and then you resume the job with Resume With Options > Replicate , the job fails. (October 2022 release)
DBMI-9234	In database ingestion initial load jobs that have an Oracle source and Snowflake target, some subtasks might fail with Oracle JDBC driver errors when the Secure Agent runs on Windows. (January 2023 release)
DBMI-9011	In a database ingestion task with an Oracle source, if you define custom data type mappings for NCHAR or NVARCHAR2 source data types with a specific length, the mapped target data type has the incorrect length in the target table. (November 2022 release)
DBMI-8809	For database ingestion incremental load and combined initial and incremental load tasks, the CDC script that the task wizard generates does not run the procedure required for enabling CDC on the Amazon Relational Database Service (RDS) for SQL Server source database. (October 2022 release)
DBMI-7753	While a database ingestion job is running, its connection might be dropped because of a problem with the Hikari pooling mechanism. (November 2022 release)
DBMI-7075	After you test an import of exported assets on the Import Assets page, you cannot change the selection of assets that failed the test because the selection check boxes are disabled. Consequently, you cannot revise the set of selected assets based on the test results. (October 2022 release)

Known issues

The following table describes Mass Ingestion Databases known issues. Note that the (*month year*) value in each issue description represents the release in which the issue was found.

CR	Description
DBMI-12203	<p>A database ingestion initial load job with a SQL Server source that includes a <code>sql_variant</code> column and a SQL Server target converts the <code>sql_variant</code> data to hexadecimal format on the target.</p> <p>Workaround: To view the original data, run the following query that uses the in-built <code>CONVERT</code> function on the target:</p> <pre>Select <column_name>, CONVERT(varbinary,<column_name>) from <table_name>;</pre> <p>Replace <code><column_name></code> and <code><table_name></code> with the actual target column and table names. (January 2023 release)</p>
DBMI-12115	<p>Database ingestion initial load jobs that have a SQL Server target and process source data that contains special characters <code>"\n"</code> or <code>"\r"</code> might end with the following error:</p> <pre>Error:Unable to retrieve data from the source com.microsoft.sqlserver.jdbc.SQLServerException: Failed to parse the CSV file, verify that the fields are correctly enclosed in double quotes.</pre> <p>Workaround: None. (January 2023 release)</p>
DBMI-12112	<p>Database ingestion incremental load jobs that have a SQL Server source and Snowflake target and that use the query-based CDC technique might fail when the time switches between Daylight Savings Time (DST) and standard time. The failure occurs if the time switch occurs while the job is running.</p> <p>Workaround: Undeploy the job and then update the initial start time for the task to account for the DST offset. Then redeploy the job to put the updated start time into effect. (January 2023 release)</p>
DBMI-11987	<p>A database ingestion initial load job with a SQL Server source table and a SQL Server target might fail with a <code>java.lang.StackOverflowError</code> error when reading a large source data record with multiple special characters, especially double quotation marks (<code>"</code>).</p> <p>Workaround: None. (January 2023 release)</p>
DBMI-11852	<p>When a database ingestion task that has a MySQL source with a JSON column is deployed, the target schema is generated without the corresponding target column. A target column should be present even though JSON source columns are not supported and only nulls would be replicated for these columns.</p> <p>Workaround: None. (January 2023 release)</p>
DBMI-11732	<p>If database ingestion incremental load or combined initial and incremental load jobs replicate LOB source data to an Amazon S3, Google Cloud Storage, or Microsoft Azure Data Lake Storage Gen2 target and use the CSV format for the target output file, the LOB data appears as empty strings in the target file.</p> <p>Workaround: None. (November 2022 release)</p>
DBMI-11681	<p>When a database ingestion task that has a MySQL source with a JSON column is deployed, the target schema is generated without the corresponding target column and the column does not appear on the target. A target column should be present even though JSON source columns are not supported and only nulls are replicated for them.</p> <p>Workaround: None. (January 2023 release)</p>

CR	Description
DBMI-11611	<p>If you set the SQL Endpoint JDBC URL property in the connection properties for a Databricks Delta target and then try to define a database ingestion task that uses the connection, the task wizard cannot retrieve the list of target schemas.</p> <p>Workaround: If you remove the SQL Endpoint JDBC URL value in the connection properties, the task wizard can use the connection to retrieve the schemas. (October 2022 release)</p>
DBMI-11552	<p>Database ingestion initial load jobs that use the Informatica-supplied Progress DataDirect JDBC driver for SQL Server to connect to a SQL Server source fail.</p> <p>Workaround: Download the Microsoft JDBC Driver for SQL Server. (January 2023 release)</p>
DBMI-11515	<p>The deployment of a database ingestion job that has an SAP HANA source and Google BigQuery target might fail if some SAP HANA table names include a forward slash (/) or a hyphen (-) and you define rules to rename the tables on the target. The forward slash and hyphen characters are replaced by underscores (_) in the target table names.</p> <p>Workaround: None. (September 2022 release)</p>
DBMI-11422	<p>Database ingestion incremental load jobs that have one or more MySQL source columns with the SET or ENUM data type and a Snowflake target replicate the SET and ENUM column data as numeric values to the target.</p> <p>Workaround: None. (October 2022 release)</p>
DBMI-11248	<p>If you try to share a Secure Agent between a parent organization and suborganization, when you create an application ingestion task or database ingestion task under the parent organization or suborganization, the task wizard displays a message that warns that you do not have sufficient permissions to access the runtime environment with the shared Secure Agent. You cannot create the task.</p> <p>Workaround: None. (October 2022 release)</p>
DBMI-11218	<p>If you set any schema drift option to Stop Table for a database ingestion job that has a PostgreSQL source, when a DDL change of that type occurs on the source, the job fails. When you try to resume the job with the Resume with Options > Replicate option, the resumed job fails to capture the DDL changes and does not issue an alert notification.</p> <p>Workaround: None. (August 2022 release)</p>
DBMI-11186	<p>If you define a Microsoft Azure Synapse Analytics Database Ingestion target connection with the setting Authentication=ActiveDirectoryPassword in the Azure Synapse Analytics JDBC URL field and then test the connection, the test fails.</p> <p>Workaround: None. You can ignore the test failure and still be able to create and run the task with the connection. (January 2023 release)</p>
DBMI-11034	<p>Database ingestion initial load jobs that have an Oracle target fail if SSL use is enabled in the Oracle Database Ingestion connection properties.</p> <p>Workaround: None. (August 2022 release)</p>
DBMI-11033	<p>If you specify custom data type mapping rules for a Snowflake target in a database ingestion task that has an SAP HANA source, the mapping rules are not applied. The target schema that is generated when you deploy the task does not include the expected target data types.</p> <p>Workaround: None. (August 2022 release)</p>
DBMI-10794	<p>Oracle source columns with the TIMESTAMP WITH TIME ZONE data type are supported only for initial load jobs. (July 2022 release)</p>

CR	Description
DBMI-10342	<p>Database ingestion incremental load and combined load jobs that have a SQL Server source might fail if SSL encryption is enabled in the connection properties. This problem occurs because the trust store file cannot be loaded even though it is specified in the connection properties. The following error is issued:</p> <pre>PWX-36749 Program internal error. Prog="SSLR". Line=0. Desc="[1] ODBC driver returned error 28000 [Informatica][ODBC SQL Server Wire Protocol driver]Cannot load trust store. SSL Error Queue:(0)".</pre> <p>Workaround: In the connection properties, set the Validate Server Certificate option to false and remove the Trust Store value. Then try to run the job again. (July 2022 release)</p>
DBMI-10210	<p>Subtasks of database ingestion initial load jobs that have a PostgreSQL source and Snowflake target might fail if source table LOB (JSON or XML) columns contain more than 16777216 bytes of data and are mapped to Snowflake variant columns. This mapping is the default data type mapping.</p> <p>Workaround: Create a custom data type mapping that maps the LOB source columns to target columns that have a data type other than variant, such as varchar. Or set the unloadClobTruncationSize custom property to a value less than 16777216 to reduce the truncation size. (July 2022 release)</p>
DBMI-10193	<p>If you create a custom data type mapping that maps VARCHAR to NVARCHAR(MAX) in a database ingestion task with a Db2 for LUW source and Azure Synapse Analytics target, after you run the associated job, the target column will have the VARCHAR data type instead of the expected NVARCHAR(MAX) data type. (July 2022 release)</p>
DBMI-9852	<p>If you download and manually run the CDC script for a database ingestion task that has a SAP HANA source, the PKLOG and PROCESSED tables might be created in a schema that is different from the schema that is specified for those tables in the body of the trigger DDL, if you specified an override schema in the advanced connection properties or the current user who is executing the script has a different default schema.</p> <p>Workaround: Modify the downloaded script to ensure that the PKLOG and PROCESSED tables are created in the schema that is specified for them within the body of the CREATE TRIGGER DDL. (May 2022 release)</p>
DBMI-9823	<p>Database ingestion incremental load jobs that have a SQL Server source with table columns defined with the sql_variant data type fail when processing the first DML operation on the table.</p> <p>Workaround: Remove the table from the task or change the data type of the sql_variant columns. (May 2022 release)</p>
DBMI-7483	<p>Deployment of database ingestion initial load tasks that have a SAP HANA source and Microsoft Azure Synapse Analytics target might fail if a source table contains a multiple-column primary key of a long length.</p> <p>Workaround: Reduce the length of the primary key. (May 2022 release)</p>
DBMI-7480	<p>A database ingestion incremental load or combined initial and incremental load job that has an Aurora PostgreSQL source and Google BigQuery target fails after an Add Column change when the schema drift option is Stop.</p> <p>Workaround: Change the schema drift option to Replicate and redeploy the job. (December 2021 release)</p>

CR	Description
DBMI-7443	<p>An attempt to deploy a database ingestion initial load task that has a SAP HANA source and a Microsoft Azure Synapse Analytics target fails with a Metadata Handler error after the target table has been created if the source table contains a column that has the unsupported data type of TEXT or SHORTTEXT. Previously, the deployment succeeded and the job instance wrote nulls to the target for these columns.</p> <p>Workaround: Try deploying the task again. If the problem still exists, change the data type of the source column, drop the column, or edit the task to exclude the table that contains the column. (November 2021 release)</p>
DBMI-7374, DBMI-7376, DBMI-7380	<p>If you deploy an initial load task that has an Amazon Relational Database Service (RDS) for MySQL source and an Amazon Redshift, Amazon S3, Google Cloud Storage, Microsoft Azure Data Lake Storage Gen2, Microsoft Azure Synapse Analytics, or Snowflake target, the job fails with an error message if the table columns have any of the following data types:</p> <ul style="list-style-type: none"> - bit with precision equal to or greater than 8 - tinyint(1) - year <p>Workaround: None. (November 2021 release)</p>
DBMI-7206	<p>In database ingestion tasks, source columns that have a character data type are mapped to Google Big Query string columns with no length specification.</p> <p>Workaround: If you want the target table columns to have a specific length, modify the target column definitions to specify the string column length after you deploy the task. (May 2022 release)</p>
DBMI-7204	<p>In database ingestion tasks, SAP HANA source columns that have the longdate or timestamp data type are mapped to Google Big Query columns string columns with no length specification.</p> <p>Workaround: If you want the target table columns to use a specific length, such as string(27), or to use the Google Big Query timestamp data type, modify the target table column definitions after you deploy the task. (November 2021 release)</p>
DBMI-7120	<p>If you enabled the handling of source table deletes as soft deletes on the target for database ingestion incremental load and combined initial and incremental load jobs, primary key updates will lead to data loss.</p> <p>Workaround: Do not perform an update on the primary key in a source table, or do not use the Soft Delete apply mode. (October 2021 release)</p>
DBMI-6074	<p>If you redeploy a database ingestion initial load job that is in the Completed state, the job transitions to the Up and Running state but then immediately reverts to the Completed state. In the job details, the Job Overview tab displays incorrect record counts.</p> <p>Workaround: Use the Run command to run the job again. (July 2021 release)</p>
DBMI-5373	<p>Database ingestion jobs convert infinity and NaN values in FLOAT columns in source tables to null values in Amazon Redshift target tables. (May 2021 release)</p>
DBMI-5340	<p>Unload subtasks for a database ingestion job with an Oracle source and flat file target might randomly fail with the following error:</p> <pre>null value in entry: unloadDatabasePassword=null</pre> <p>Workaround: Run the job again, or create a new one. (May 2021 release)</p>
DBMI-5311	<p>The deployment of tasks that have an Amazon Redshift target fails if the name of a source table or column has the maximum length of 128 characters. (May 2021 release)</p>

CR	Description
DBMI-5218	<p>Database ingestion jobs that load data to Amazon Redshift targets fail under the following conditions:</p> <ul style="list-style-type: none"> - The source table contains columns that have a string data type and contain data with more than 16383 UTF-8 characters or 65535 ASCII characters. - The source table contains columns that have a binary data type and contain data with more than 32767 characters. <p>(April 2021 release)</p>
DBMI-5087	<p>For database ingestion incremental load and combined initial and incremental load tasks that have an Oracle source, the supplemental logging settings for tables might be ignored by Oracle if the table name or any table column name is longer than 30 characters. In this case, the results are unpredictable.</p> <p>Workaround: Avoid replicating tables that have a table or column name longer than 30 characters. (March 2021 release)</p>
DBMI-4239	<p>Incremental database ingestion jobs end with an error if a source schema or table name includes an asterisk (*), percent character (%), or a question mark (?). (February 2021 release)</p>
DBMI-3711	<p>When you open multiple database ingestion job or task details and switch between them on the left panel, Mass Ingestion Databases does not refresh job statistics and task information. (November 2020 release)</p>
DBMI-3652	<p>Mass Ingestion Databases might fail to deploy a task that has a Microsoft SQL Server source with the error 'SSL peer shut down incorrectly'.</p> <p>Workaround: Set the SQL Server Keep Alive option for the SQL Server source to 20 seconds or greater. (November 2020 release)</p>
DBMI-3631	<p>Database ingestion incremental load jobs that have a Microsoft Azure Synapse Analytics target might fail when trying to create an external table if a source table has more than 508 columns.</p> <p>Workaround: Make sure that the selected tables have no more than 508 columns and the columns have supported data types. (November 2020 release)</p>
DBMI-2957	<p>For Oracle sources, when a new column with a default value is added to an existing source table or an existing column is modified to have a default value, Mass Ingestion Databases does not propagate the default value to any existing rows on the target. This issue occurs even though the Add Column schema drift option is set to Replicate. If Updates occur on the source, duplicate rows might be written to the target if the source table does not have a primary key. Also, deletions might not be processed if the source table does not have a primary key. (July 2020 release)</p>
DBMI-2885	<p>For Microsoft Azure Synapse Analytics targets, if you change a secret key for the Microsoft Azure Data Lake Storage (ADLS) Gen2 staging area and update the Client Secret value in your connection properties, incremental load and combined initial and incremental load jobs that use the staging area end with an error.</p> <p>Workaround: After you change the secret key, use the following commands to drop the staging file and database credential on the ADLS Gen2 account:</p> <pre>DROP EXTERNAL DATA SOURCE data_source; DROP DATABASE SCOPED CREDENTIAL credential;</pre> <p>Note: You might also have to drop the external tables that use the external data source. (July 2020 release)</p>

CR	Description
DBMI-2875	If you rename a primary key column for a source table in an incremental load job or in a combined initial and incremental load job, Mass Ingestion Databases incorrectly treats the change as either a dropped primary key or as a column dropped from a composite primary key. In both cases, Mass Ingestion Databases issues an alert message and stops processing the table. Workaround: Resynchronize the table. (July 2020 release)
DBMI-2783	The Resume With Options command is not available for jobs in the Running with Warning state if one or more of the tables were excluded from replication because of the Stop Table schema drift option. Workaround: Stop the job and then click Resume with Options in the Actions menu to resume the job. (July 2020 release)
DBMI-2752	A database ingestion job might fail to resume processing a source table with the following error: [36202] NULL value found for column that is not nullable. This error occurs if all of the following events occur: <ul style="list-style-type: none"> - An ALTER COLUMN operation on a source table adds the NOT NULL constraint. - The job is configured to stop if a DDL operation occurs on a source table. - The database ingestion job tries to resume processing from the transaction that occurred just prior to the DDL operation. Workaround: Resynchronize the source table with the target. (July 2020 release)
DBMI-2673	Database ingestion incremental load and combined initial and incremental load jobs incorrectly handle the situation where a source table is dropped and then a new table with the same name is created. If the new table is identical to the dropped table, the jobs might fail to detect the change. If the old and new tables differ, the jobs handle the change as a schema drift change. In both cases, the jobs do not truncate the target table after the source table is re-created, which results in data inconsistencies. Workaround: Resynchronize the re-created source table with the target. (June 2020 release)
DBMI-2537	If you add or remove a primary key or unique constraint from a source table in an incremental load job or in a combined initial and incremental load job, Mass Ingestion Databases does not replicate the change to the target. (April 2020 release)
DBMI-1760	For Oracle sources, database ingestion jobs that perform incremental loads do not propagate Update operations on one or more pairs of rows if the Update changes the primary key value of the first row to the primary key value of the second row and the primary key value of the second row to the primary key value of the first row. (November 2019 release)
DBMI-1570	If a database ingestion task includes more than 1,000 source tables, downloading a list of the source tables that match the selection rules might fail with a gateway timeout error. (October 2019 release)

Informatica Global Customer Support

You can contact a Customer Support Center by telephone or online.

For online support, click **Submit Support Request** in Informatica Intelligent Cloud Services. You can also use Online Support to log a case. Online Support requires a login. You can request a login at <https://network.informatica.com/welcome>.

The telephone numbers for Informatica Global Customer Support are available from the Informatica web site at <https://www.informatica.com/services-and-training/support-services/contact-us.html>.