



Informatica Data Ingestion and Replication July 2024 Data Ingestion and Replication Release Notes

© Copyright Informatica LLC 2019, 2024

Contents

| | |
|--|----|
| Common. | 1 |
| Fixed issues. | 1 |
| Known issues. | 3 |
| Application Ingestion and Replication. | 3 |
| Fixed issues. | 3 |
| Known issues. | 4 |
| Database Ingestion and Replication. | 4 |
| Fixed issues. | 4 |
| Known issues. | 11 |
| File Ingestion and Replication. | 13 |
| Fixed issues. | 13 |
| Known issues. | 14 |
| Streaming Ingestion and Replication. | 14 |
| Fixed issues. | 14 |
| Known issues. | 14 |
| Informatica Global Customer Support. | 14 |

Read the *Data Ingestion and Replication Release Notes* to learn about fixed issues and known issues in the July 2024 release of the Informatica Intelligent Cloud ServicesSM Data Ingestion and Replication service, which includes Application Ingestion and Replication, Database Ingestion and Replication, File Ingestion and Replication, and Streaming Ingestion and Replication.

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

Common

Fixed issues

The following tables describe fixed issues that apply to two or more ingestion types. Not all monthly releases include fixed issues.

July 2024 fixed issues

| Issue | Description |
|------------|---|
| AIN-6535 | Application ingestion and replication and database ingestion and replication jobs that use the Superpipe option to stream data to a Snowflake target fail with the following exception when there is a large number of target tables and no LIMIT on the show columns command. SnowflakeSQLException: Statement exhausted execution resources and was canceled. Please contact support. |
| AIN-6527 | Application ingestion and replication and database ingestion and replication jobs fail if the source table name contains a forward slash (/). |
| DBMI-18843 | Application ingestion and replication jobs and database ingestion and replication jobs that have a PostgreSQL target fail with the following error if you enabled the SSL encryption method in the PostgreSQL connection properties but didn't specify a KeyStore value: Process <UNLOAD-SALESFORCE-Account PostgreSQLFile Writer> failed to initialize due to an unexpected exception creating of POSTGRESQL connection. Error : [informatica][PostgreSQL JDBC Driver]Error establishing socket to host and port: <host>:<port>. Reason: null (No such file or directory) |
| DBMI-18827 | Application ingestion and replication jobs and database ingestion and replication jobs that use the Superpipe option to stream rows of data directly to Snowflake target tables need to prevent the stream from becoming stale during periods of no change data activity, when the pointer for the last read is not advanced. If the stream becomes stale, views of the CDC target tables might not be available. |

May 2024 fixed issues

| Issue | Description |
|------------|--|
| DBMI-17836 | Application ingestion and replication and database ingestion and replication incremental load and combined initial and incremental load jobs that have a Databricks Delta target, have no primary key in the source table and use a custom setting for column name case transformation might fail with a MERGE INTO query, resulting in an improper column matching. |

April 2024 fixed issues

| Issue | Description |
|------------|---|
| DBMI-15990 | Application ingestion and replication and database ingestion and replication incremental load and combined initial and incremental load jobs that have a SQL Server source table with LOB columns but no primary key might apply duplicate records to the corresponding target table. |
| DBMI-15836 | Application ingestion and replication and database ingestion and replication incremental load and combined load jobs that have a Databricks Delta target and use Microsoft Azure Data Lake Storage (ADLS) Gen2 to stage data do not delete the parent staging directory after all CDC cycles have ended and all cycle-specific subdirectories have been deleted. The parent directory should also be deleted. |

Known issues

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

| Issue | Description |
|---------------------------|---|
| DBMI-18254 | In Data Integration - Free and PayGo, the New Asset dialog box contains a Data Ingestion and Replication option even though Data Ingestion and Replication tasks are not supported for this product. If you click the option, no ingestion and replication task types are displayed for you to select. Workaround: None. (May 2024) |
| DBMI-18190 | From the Data Integration Home page, you can set a primary cloud data warehouse destination. However, this feature is currently not supported for configuring a target for Data Ingestion and Replication tasks. Workaround: Define the target cloud data warehouse connection in Administrator and then select the connection in the task wizard when you create an ingestion and replication task. (May 2024) |
| DBMI-12846, DBMI-12847 | If you use a Secure Agent on Windows that is configured for NTLM proxy authentication, when you try to create an application ingestion and replication task or database ingestion and replication task that has an Amazon RDS for Oracle or Amazon Redshift target, the task wizard fails to fetch the list of target schemas. As a result, you cannot complete the task definition. This problem does not occur with an on-premises Oracle target. Workaround: None. (April 2023 release) |

Application Ingestion and Replication

Fixed issues

The following tables describe recent Application Ingestion and Replication fixed issues. Not all monthly releases include fixed issues.

April 2024 fixed issues

| Issue | Description |
|----------|---|
| AIN-2981 | Application ingestion and replication incremental load or combined initial and incremental load jobs that have Snowflake targets and use the Apply Mode setting of Soft Delete replicate Update operations on the source as a pair of Delete and Insert records, which result in Delete records appearing on the target when Update records are expected. |

February 2024 fixed issues

| Issue | Description |
|----------|--|
| AIN-4755 | If you delete data from a child table in an application ingestion and replication incremental load or combined load task with a ServiceNow source, the data is also deleted from the parent table. This applies only if the Secure Agent properties <code>metadata_manager.jvm.ai.servicenow.enable.cdc.using.inherited.columns</code> and <code>task_container.jvm.ai.servicenow.enable.cdc.using.inherited.columns</code> are set to true. |

Known issues

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

| Issue | Description |
|----------|--|
| AIN-4591 | Application ingestion and replication incremental load and combined load jobs with an Oracle Fusion BICC source don't create a subtask if you use a dot (.) as a datastore separator. Workaround: Use an asterisk (*) instead of dot (.). (August 2023) |

Database Ingestion and Replication

Fixed issues

The following tables describe recent Database Ingestion and Replication fixed issues. Not all monthly releases include fixed issues.

July 2024 fixed issues

| Issue | Description |
|------------|---|
| DBMI-19217 | If you select the advanced Add Last Replicated Time property on the Target page of the task wizard and save task and then later clear the check box and save the task again, when you return to the task wizard, the check box is still selected. However, if you open the task in view mode, the property is disabled. In this case, the database ingestion and replication job does not add the last replicated time metadata column on the target, which appears to conflict with the wizard setting. |
| DBMI-18809 | Database ingestion and replication jobs that have a Snowflake target and use the Superpipe feature to stream change data to target tables might fail with the following error if mixed-case and lowercase column identifiers are quoted and if the <code>QUOTED_IDENTIFIERS_IGNORE_CASE=true</code> parameter is set in the Additional JDBC URL Parameters field of the connection properties: Exception handled: 'SQL compilation error: error line 1 at position <bytes> invalid identifier '<identifier>''. With error code 904 |

| Issue | Description |
|------------|--|
| DBMI-18786 | <p>In the Data Ingestion and Replication Command-Line Interface (CLI), if you issue the <code>job status</code> command without the optional <code>-f</code> option, the command fails with the following error:</p> <pre>Could not get the fields for job '<job_instance>' in the 'Up and Running' state.</pre> <p>The failure occurs because of date-time format discrepancies that are encountered when fetching the job's start time.</p> |
| DBMI-18774 | <p>Database ingestion and replication jobs that have an Amazon Redshift target might fail when creating target tables that have names in mixed case or uppercase. The failure occurs if the Amazon Redshift <code>enable_case_sensitive_identifier</code> cluster configuration parameter is set to false and the Enable Case Transformation target property is not selected for the job.</p> |
| DBMI-18670 | <p>Database ingestion and replication jobs that use a Microsoft Azure Data Lake Storage Gen2 target connection to connect to an Azure Government cloud fail because the connection uses a specific host name by default that is incorrect. With this fix, the connection to the Azure Government cloud succeeds provided that you use Azure Service Principal authentication.</p> |
| DBMI-18650 | <p>If you create and deploy an application ingestion and replication task or database ingestion and replication task that has a database target in the Data Ingestion and Replication CLI, the task is created without an apply mode setting and the default Standard apply mode is not used. Consequently, when you try to run the associated job instance, it fails.</p> |
| DBMI-18629 | <p>In a database ingestion and replication task that has an Oracle source and Google BigQuery target, if you add a custom data type mapping that maps the Oracle DATE type to the Google BigQuery TIMESTAMP type, the custom mapping is ignored when you run the job. Instead, the default mapping of Oracle DATE to Google BigQuery DATETIME is used.</p> |
| DBMI-18243 | <p>If you deselect the Add Operation Type metadata column for database ingestion and replication incremental load or combined load tasks that have an SAP HANA source and Snowflake target and that use Audit apply mode, the following exception occurs when you try to run the associated job:</p> <pre>java.lang.IndexOutOfBoundsException: Index: 139, Size: 139</pre> |
| DBMI-18233 | <p>Database ingestion and replication incremental load jobs that have a Db2 for i source fail repeatedly before issuing the following error message:</p> <pre>[CDCPUB_10066] TRACE: [IBMiLogCollector failed creating SQL IBMiLogCollector createNewLogSQLStatement(), current log start sequence <sequence_number> is greater than ending log sequence <sequence_number> for the journal receiver being processed.]</pre> |
| DBMI-18230 | <p>In database ingestion and replication tasks that have an Oracle source, selected tables might be automatically deselected in the following situation:</p> <ol style="list-style-type: none"> 1. You apply an Include * rule that includes all tables in the schema. 2. Under Table View, you deselect all tables. 3. Under Table View, you search for a table and then select a subset of its columns. 4. You repeat step 3 multiple times. |
| DBMI-18216 | <p>During database ingestion and replication CDC processing, the performance of queries to the Db2 for i journal receivers for data needs to be improved.</p> |
| DBMI-18169 | <p>A database ingestion and replication job that has a Kafka target configured with one-way SSL connection mode does not start because of a missing SSL KeyStore Password connection property which is not required for this mode.</p> |

| Issue | Description |
|------------|---|
| DBMI-18098 | A database ingestion and replication initial load job that has a Microsoft Azure Data Lake Storage Gen2 target does not attempt a retry if the upload to target fails. |
| DBMI-18097 | <p>Database ingestion and replication jobs that have a SQL Server source and SQL Server target might fail with the following error when processing large queries to the database:</p> <p>Caused by: java.sql.SQLException: [informatica] [SQLServer JDBC Driver] [SQLServer] There is already an object named '<temp_table_name>_VER' in the database.</p> <p>The problem persists even after the job is undeployed, the tables on the target are dropped, and the job is deployed again.</p> |
| DBMI-16778 | <p>Redeployed database ingestion and replication combined load jobs that have an Oracle 11g source might fail with the following error when querying redo logs for start information:</p> <p>PWX-36075 OCI Error: ORA-00907: missing right parenthesis.</p> |
| DBMI-13489 | Database ingestion and replication incremental load and combined load jobs that have a Db2 for i source need a way to prevent the deletion of source journal receivers during CDC processing. |

May 2024 fixed issues

| Issue | Description |
|------------|--|
| DBMI-18099 | Database ingestion and replication incremental load or combined initial and incremental load jobs that have Db2 for i sources might fail if the query for source journal receivers encounters an empty journal receiver name. |
| DBMI-18096 | Database ingestion and replication incremental load jobs that have Oracle sources might fail if the TDE master key is changed. |
| DBMI-17805 | Database ingestion and replication incremental load jobs that have Oracle sources and run in ARCHIVECOPY reader mode end abnormally after a failover from the primary database to a standby database. |
| DBMI-17593 | The Resync option is available for a database ingestion and replication job even though the user does not have the required privileges to perform a resync. |
| DBMI-17552 | During database ingestion and replication CDC processing of an Oracle source, a capture progress marker might be sent out of sequence after processing catches up to the current end of log. |
| DBMI-17547 | In database ingestion and replication jobs that have an Oracle source and a Snowflake target with schema drift enabled, the target column size might be inconsistent with the source column size. |
| DBMI-17202 | When running a database ingestion and replication initial load job that has a MySQL source, the retrieval of the first record might be delayed by hours, reducing the overall throughput of the database ingestion job. In this case, an out-of-memory issue might occur if the source table result set is large enough to exceed the available heap memory of the database ingestion and replication task container process. In combined initial and incremental load jobs, the switch to incremental load processing might also be delayed by hours. |

| Issue | Description |
|------------|--|
| DBMI-17189 | Database ingestion and replication incremental load or combined load jobs that use the Soft Deletes apply mode might not apply all change records to the Snowflake target if you delete a source record and then add the same record shortly afterwards again. In this case, only the Delete record is captured. |
| DBMI-17181 | Database ingestion and replication combined initial and incremental load jobs that have a MySQL source might fail with the following error: ADAPTER MySQL <Client> Not able to get end of log position. Failed with error: The last packet successfully received is longer than the server configured value of 'wait_timeout' |
| DBMI-17179 | When running a database ingestion and replication initial and incremental load job, the switch from initial load processing to incremental load processing might be delayed by hours. |
| DBMI-17174 | Database ingestion and replication incremental load jobs that have a Db2 for z/OS source fail repeatedly with the following error message: [CDCPUB_10066] TRACE: [PwxCDCReaderRunHandler encountered error :Previous sequence <sequence_number> is not less than current sequence <sequence_number>. Caused by: Previous sequence <sequence_number> is not less than current sequence <sequence_number> |
| DBMI-17019 | Database ingestion and replication jobs might corrupt multi-byte character data in SAP HANA source tables when replicating the data to Snowflake targets. The problem occurs with encodings other than UTF-8. |
| DBMI-16578 | Database ingestion and replication incremental load or combined initial and incremental load jobs that use a Microsoft Azure Synapse Analytics Database Ingestion connection to the target might fail with a target connection reset error during the execution of merge apply statements at the target database. |
| DBMI-16441 | Database ingestion and replication jobs that have a Snowflake target and the Persistent Storage and Superpipe options enabled might fail when a new internal StreamObjectID is assigned to the source object schema. |
| DBMI-16384 | Database ingestion and replication jobs that have an Oracle source might ignore transactions with no DML activity of interest after 60 minutes. Data loss might occur after a job restarts if an ignored transaction that's no longer being tracked later contains DML changes. |
| DBMI-16170 | Database ingestion and replication incremental load jobs now detect non-responsive connections and fail if a connection is non-responsive for 5 minutes. This situation is usually caused by a network failure that leaves dead sockets. |
| DBMI-16138 | Database ingestion and replication combined initial and incremental load jobs that have an Oracle source might incorrectly report the suspension time in the suspension duration trace message. |
| DBMI-15421 | When database ingestion and replication jobs read Oracle logs using BFILE trailing directory separators in the Oracle directory object, the read operation might fail. |
| DBMI-11880 | If you Abort a running database ingestion and replication initial load job for which partitioning is enabled, the job does not end immediately. Abort execution is delayed. |

April 2024 fixed issues

| Issue | Description |
|------------|---|
| DBMI-16940 | <p>Database ingestion and replication incremental load jobs that have a Db2 for z/OS source and Snowflake target might fail repeatedly with the following error:</p> <pre>CDCPUB_10066] TRACE: [PwxCDCReaderRunHandler encountered error :Previous sequence <previous_sequence_number> is not less than current sequence <current_sequence_number>. Caused by: Previous sequence <sequence_number> is not less than current sequence.</pre> |
| DBMI-16392 | <p>A database ingestion and replication job that has a Databricks Delta target might fail if the source table contains NULL values in a Timestamp column. The problem occurs when the Databricks Delta ANSI_MODE configuration parameter is set to TRUE and the job tries to use the Databricks Delta MERGE statement with the COALESCE function to process NULL values for the column.</p> |
| DBMI-16382 | <p>During periods of no DML activity, database ingestion and replication combined initial and incremental load jobs that have a Db2 for i source might hang with subtasks in the Queued state during the initial load phase. The job log contains the following types of messages:</p> <pre>[CDC_COMBINED Transition Scheduler Helper(264)] INFO com.informatica.msglogger - [DBMIP_23019] The process [CDC_COMBINED] with the thread ID [264] is waiting for in-flight data from a previous run to complete. Minutes waited [17]. [CDC_COMBINED Transition Scheduler Helper(264)] INFO com.informatica.msglogger - [DBMIP_23020] Thread [264] [PUBLISHED] transition event [SCHEDULER_REQUEST_INFLIGHT_DATA_MARKER]. Details: [, time=1708074372616].</pre> |
| DBMI-16290 | <p>Database ingestion and replication combined initial and incremental load jobs that have an Oracle source and Snowflake target and schema drift enabled might fail if you increase the precision of a column, for example, from VARCHAR(16) to VARCHAR(20). The problem occurs if you use code page CP1252 at task deployment time, but this character set is unavailable at runtime. At runtime, the job will use the UTF-8 code page, which causes the precision to be reduced and the following exception:</p> <pre>[Thread-156177] LEVEL_1 trace - Exception handled: 'SQL compilation error: cannot change column <column_name> from type VARCHAR(20) to VARCHAR(16) because reducing the byte-length of a varchar is not supported. With error code 40050</pre> |
| DBMI-16277 | <p>A means to override the code page used with a coded character set identifier (CCSID) encoding is not available. With this fix, you can use the <code>pwx.cdcreader.iseries.option.CCSID2CodePageOverrides</code> (<i>ccsid,code_page</i>) custom property to specify an override code page for a given CCSID at the Db2 for i table level. For more information, contact Informatica Global Customer Support.</p> |

| Issue | Description |
|------------|--|
| DBMI-16257 | <p>Database ingestion and replication jobs that have Db2 for i sources cannot retry queries for getting the journal receiver list from Db2 for certain SQL error codes, such as:</p> <pre>[CDCPUB_10066] TRACE: [IBMiClient setErrorCondition() called. Error<IBMiMetadataManager getSchemas(), Exception JDBCapi ExecutesQL(), -443 58004 [informatica][DB2 JDBC Driver][DB2]An error occurred executing function or procedure QSYS.QDBSSUDF2 with the following parameters: (SYSTEM ERROR CONDITION 2). Refer to the joblog for more information regarding the detected error.SQL Command: ...</pre> <p>With this fix, you can use the <code>pxw.cdcreader.iseries.option.useNewLogCollector=true</code> and <code>pxw.cdcreader.iseries.option.useJournalReceiverQueries= true</code> custom properties with the <code>pxw.cdcreader.iseries.option.ConnectionRetryCodes</code>, <code>pxw.cdcreader.iseries.option.journalReceiverRetries</code>, and <code>pxw.cdcreader.iseries.option.journalReceiverInterval</code> custom properties to enable retry attempts when encountering SQL errors during journal receiver queries. For more information, contact Informatica Global Customer Support.</p> |
| DBMI-16199 | The <code>pxw.cdcreader.ZOS.ConnectionRetryCodes</code> custom property should accept negative error codes such as -443. |
| DBMI-16188 | <p>Database ingestion and replication incremental load jobs that replicate data from an Oracle 19c source with LOB columns and with non-LOB columns defined with the NOT NULL DEFAULT constraint might fail with following message when processing some Insert operations:</p> <pre>PWX-36200 ORAD: Column conversion error: [36202] NULL value found for column that is not nullable for Table name=<table_name>, column[<number>] segment[<number>] ADULT_INDICATOR. Detail Type = VARCHAR, Oracle Type = CHAR.</pre> |
| DBMI-16186 | When database ingestion and replication jobs query Db2 for i source journal receivers, the jobs might fail with SQLCODE -443 if the SQL query is created with an invalid sequence number. |
| DBMI-16185 | When you run a database ingestion and replication job that has a Db2 for z/OS source, the Db2 for z/OS log parser might not detect a failure in schema drift parsing, causing the job to continue to run without capturing any data. |
| DBMI-16184 | <p>Database ingestion and replication incremental load jobs that have an Oracle source with some LOB columns might fail with the following message when trying to process some Update operations:</p> <pre>PWX-36045 ORAD WarnOPASM: OP:26.1 unexpected unexpected operation type found at offset <bytes>. pKDLI->GetOp()= 2, length= <bytes> [PwxOrlCmnCV:2643] at redo log position SCN <scn> RBA <rba> (file: <number>).</pre> <p>The problem occurs if Oracle generates a bad redo record for the Update.</p> |
| DBMI-16049 | After you save a database ingestion and replication task that has Oracle or SQL Server source tables that are selected based on rules, some with a subset of selected columns, if you then add or remove selected columns for some tables and click the Refresh icon, the Table View count of source tables and the column selection for some tables are incorrect. |
| DBMI-16014 | Database ingestion and replication incremental load jobs that have an on-premises SQL Server source and an Azure SQL Database target incorrectly replicate Update operations to the target when multiple Updates occur in a transaction for the same column using the same primary key. |

| Issue | Description |
|------------|--|
| DBMI-16013 | Database ingestion and replication jobs that have a Kafka target might log sensitive information, such as a password, in clear text in the job log if it was entered in the Additional Connection Properties field of the connection properties. |
| DBMI-15796 | Database ingestion and replication jobs that have source columns with the FLOAT data type and precision greater than 32 and that use the Parquet output format do not convert the source FLOAT columns to target data types consistently across all load types. Initial load jobs generate target FLOAT columns, whereas incremental load and combined load jobs generate target DOUBLE columns. Now, jobs of all load types generate target DOUBLE columns. |

February 2024 fixed issues

| Issue | Description |
|------------|--|
| DBMI-16104 | A database ingestion and replication job that has a SQL Server source and uses the query-based CDC method might incorrectly replicate INTEGER data with the value of NULL to the target as 0 and replicate BIT data with the value of NULL as FALSE. |
| DBMI-15903 | A database ingestion and replication job that has a Db2 for i source might not capture any change data after a journal reset in new log collection mode. |
| DBMI-15842 | When you create a new database ingestion and replication task that has an Oracle source or update an existing task, the following error might be issued when you click Apply Rules : Cannot retrieve the list of objects. HTTP response status code is 500, Duplicate key |
| DBMI-15714 | A database ingestion and replication job that has a SQL Server source and partitioning enabled might fail if you use Windows authentication to connect to SQL Server. |
| DBMI-15552 | A database ingestion and replication incremental load or combined load job that has an Oracle or SQL Server source and schema drift enabled might fail with one of the following errors, instead of issuing a schema drift alert, if you select additional source columns and redeploy the job after the job has run: A source table schema change should stop the table. Added column found :column_name conflicts with the previously ignored drop of column <same_column_name>. A change to the source definition of table [schema.table] was encountered. The table is configured to fail. The target will no longer receive data or this table. Change description(s): [Added column found :column_name conflicts with previously ignored drop of column <same_column_name>]. |
| DBMI-15527 | Database ingestion and replication incremental load jobs that have a SAP HANA source might hang while waiting for incoming change data, although the data is present in the source tables. The job continues to run without reporting any errors. |
| DBMI-15494 | On the Source page of the database ingestion and replication task wizard, the Table View shows only previously selected columns for a table, instead of all selected and unselected columns, if you filtered for that table in the table view after the task was in the Deployed, Completed, or Up and Running state. |

| Issue | Description |
|------------|--|
| DBMI-15373 | <p>If you select the Superpipe option for a database ingestion and replication job that has a SQL Server source and a Snowflake Data Cloud target and if a source column has the DATETIMEOFFSET data type, the job fails with the following error when trying to write data for that source column to a TIMESTAMP column on the target:</p> <pre>java.lang.RuntimeException: net.snowflake.ingest.utils.SFException: The given row cannot be converted to the internal format due to invalid value: Value cannot be ingested into Snowflake column "column_name" of type TIMESTAMP, Row Index: 0, reason: Not a valid value,...</pre> |
| DBMI-15321 | A database ingestion and replication job that a SQL Server source and uses the query-based CDC method might fail to load the source tables if you select Windows Authentication v2 as the Authentication Mode in the SQL Server connection properties. |
| DBMI-14797 | A database ingestion and replication task might fail with a SQL compilation error if Database Ingestion and Replication detects a Drop Column request for a column that was previously removed. |
| DBMI-13202 | <p>Database ingestion and replication combined initial and incremental load jobs that have an Oracle source with a VARCHAR2 column and an Oracle target fail when incremental change data processing starts. The failure occurs because of the following unique constraint error:</p> <pre>Exception handled: [informatica][Oracle JDBC Driver][Oracle]ORA-00001: unique constraint <constraint> violated\n'. With error code 1</pre> |

Known issues

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

| Issue | Description |
|------------|---|
| DBMI-19145 | <p>Database ingestion and replication jobs that use the Log-based CDC method and have an Oracle source with only LOB columns selected and no primary key do not create subtasks on the tables. As a result, the jobs cannot capture change data from the tables. Messages such as the following are written to the log:</p> <pre>[DBMIP_23026] The process [CDC-LOBS] with the thread ID [9708] encountered new table [lob_table_name]. The new table is excluded from processing. [CDCPUB_10066] TRACE: [PwxCDCRequestProcessor.askSchema() returned: Don't capture.].</pre> <p>Workaround: In the task, select some non-LOB source columns, in addition to the LOB columns, for replication if you want to continue using the table without a primary key.. (July 2024)</p> |
| DBMI-19126 | <p>Database ingestion and replication tasks that that are configured to not include LOBs in capture processing but have source tables with unsupported Oracle BFILE or LOB columns, incorrectly add table-level Exclude rules that specify conditions with the unsupported column names on the Source page of the task wizard.</p> <p>Workaround: None. (July 2024)</p> |

| Issue | Description |
|------------|---|
| DBMI-15910 | <p>Database ingestion and replication jobs that have a Db2 for i source and Amazon Aurora PostgreSQL target might fail when copying data from the local .csv file to the target table or to the LOG table (for an incremental load) if the default Secure Agent installation directory path contains spaces. This problem is caused by a known Progress JDBC driver for PostgreSQL issue. The driver does not preserve spaces in the directory path, causing the database ingestion job to not be able to find the .csv file.</p> <p>Workaround: Configure the Secure Agent installation directory path without spaces. (February 2024)</p> |
| DBMI-15247 | <p>Database ingestion and replication jobs that have an SAP HANA source table with a multiple-column primary key and that use a custom data-type mapping to map a source TIMESTAMP column to a Snowflake target VARCHAR column replicate data to the target incorrectly, which results in invalid data in the target column.</p> <p>Workaround: None. (November 2023)</p> |
| DBMI-14370 | <p>Database ingestion and replication incremental load jobs that have an SAP HANA Cloud source with a DECIMAL column that has no precision fail with the following error message:</p> <p>Error executing query job. Message: Query error: Value of type BIGNUMERIC cannot be assigned to column, which has type INT64 at [number]. Error code: 100032</p> <p>Workaround: None. (August 2023)</p> |
| DBMI-13605 | <p>The Oracle Database Ingestion connection properties page includes no property for entering JDBC connection properties, such as EncryptionLevel, when they're needed.</p> <p>Workaround: In the Service Name field, you can add JDBC connection properties after the Oracle SID value, using a semicolon (;) as the separator. (April 2023 release)</p> |
| DBMI-12571 | <p>Database ingestion and replication jobs with a SAP HANA source might replicate data from REAL columns of 16 or more significant digits with a loss of precision on the target, causing data corruption.</p> <p>Workaround: None. (April 2023 release)</p> |
| DBMI-12331 | <p>If you create custom data type mappings for Oracle source binary_double and binary_float columns in a database ingestion task, the custom mappings are ignored. Instead, the target table is generated using the default mappings of binary_double > float and binary_float > real. When the database ingestion and replication job runs, nulls are written to the target float and real columns.</p> <p>Workaround: None. (February 2023 release)</p> |
| DBMI-12327 | <p>For database ingestion and replication tasks that include custom data type mapping rules, the Metadata Manager log incorrectly reports nullable columns as NOT NULL and reports not-nullable columns as NULL. However, the target columns corresponding to the mapped source columns are generated with the correct nullability.</p> <p>Workaround: None. (February 2023 release)</p> |
| DBMI-11732 | <p>If database ingestion and replication 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-11552 | <p>Database ingestion and replication 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. (September 2022 release)</p> |

| Issue | Description |
|------------|--|
| DBMI-10794 | Oracle source columns with the TIMESTAMP WITH TIME ZONE data type are supported only for initial load jobs. (July 2022 release) |
| DBMI-10272 | Database ingestion and replication incremental load or combined initial and incremental load jobs that have a SQL Server source with binary, decimal, or datetimeoffset columns and an Oracle target fail if a DML update operation is followed by an insert. Workaround: None. (May 2022 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-2297 | Although the Amazon S3 connection properties allow users to specify an IAM role, you cannot use temporary security credentials generated by the AssumeRole method of the AWS Security Token Service API to authorize user access to AWS Amazon S3 resources. Workaround: None. (April 2020 release) |

File Ingestion and Replication

Fixed issues

The following table describes recent File Ingestion and Replication fixed issues. Not all monthly releases include fixed issues.

July 2024 fixed issues

| CR | Description |
|----------|---|
| FMI-2933 | If you add a file ingestion and replication task in a taskflow and click the plus (+) icon to view the file ingestion and replication task details, the task details don't appear. (May 2024 release) |

May 2024 fixed issues

| CR | Description |
|------------|--|
| CCON-94476 | If you run a file ingestion and replication job with the If File Exists option set to Append for binary files, the file content does not get appended. (February 2024 release) |

Known issues

The following table describes File Ingestion and Replication known issues. Note that the (*month, year*) value in each issue description represents the release in which the issue was reported.

| CR | Description |
|------------|--|
| CCON-94476 | If you run a file ingestion and replication job with the If File Exists option set to Append for binary files, the file content does not get appended. (February 2024 release) |

Streaming Ingestion and Replication

Fixed issues

There were no fixed issues in this release of Informatica Intelligent Cloud Services Streaming Ingestion and Replication. Not all monthly releases include fixed issues.

Known issues

The following table describes Streaming Ingestion and Replication known issues. Note that the (*month, year*) value in each issue description represents the release in which the issue was reported.

| CR | Description |
|---------|--|
| SI-4966 | When you delete an undeployed streaming ingestion and replication job, the deleted jobs are recorded in both the My Jobs page in Data Integration and the All Jobs page in Monitor. However, when you click the job to view the details, the page becomes unresponsive. (May 2024 release) |
| SI-4934 | The connection test fails for a streaming ingestion and replication task that uses the Google PubSub connector. However, the task runs successfully. (May 2024 release) |

Informatica Global Customer Support

You can contact a Global Support Center through the Informatica Network or by telephone.

To find online support resources on the Informatica Network, click **Contact Support** in the Informatica Intelligent Cloud Services Help menu to go to the **Cloud Support** page. The **Cloud Support** page includes system status information and community discussions. Log in to Informatica Network and click **Need Help** to find additional resources and to contact Informatica Global Customer Support through email.

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>.