



PowerExchange Adapters for Informatica

10.4.1

Release Notes

June 2020

© Copyright Informatica LLC 1993, 2021

Contents

PowerExchange for Amazon Redshift.....	4
PowerExchange for Amazon Redshift (10.4.1).....	4
PowerExchange for Amazon Redshift (10.4.0.1).....	5
PowerExchange for Amazon Redshift (10.4.0).....	5
PowerExchange for Amazon S3.....	8
PowerExchange for Amazon S3 (10.4.1).....	8
PowerExchange for Amazon S3 (10.4.0.1).....	9
PowerExchange for Amazon S3 (10.4.0).....	9
PowerExchange for Cassandra.....	12
PowerExchange for Cassandra (10.4.1).....	12
PowerExchange for Cassandra (10.4.0).....	12
PowerExchange for Cassandra JDBC.....	13
PowerExchange for Cassandra JDBC (10.4.1).....	13
PowerExchange for Cassandra JDBC (10.4.0).....	13
PowerExchange for DataSift.....	13
PowerExchange for DataSift (10.4.0).....	13
PowerExchange for Facebook.....	14
PowerExchange for Facebook (10.4.0).....	14
PowerExchange for Google Analytics.....	14
PowerExchange for Google Cloud Analytics (10.4.1).....	14
PowerExchange for Google Analytics (10.4.0).....	14
PowerExchange for Google BigQuery.....	15
PowerExchange for Google BigQuery (10.4.1).....	15
PowerExchange for Google BigQuery (10.4.0).....	16
PowerExchange for Google Cloud Spanner.....	17
PowerExchange for Google Cloud Spanner (10.4.1).....	17
PowerExchange for Google Cloud Spanner (10.4.0).....	17
PowerExchange for Google Cloud Storage.....	17
PowerExchange for Google Cloud Storage (10.4.1).....	17

PowerExchange for Google Cloud Storage (10.4.0).....	17
PowerExchange for Greenplum.....	18
PowerExchange for Greenplum (10.4.1).....	18
PowerExchange for Greenplum (10.4.0).....	18
PowerExchange for Greenplum Fixed Limitations (10.4.0).....	18
PowerExchange for Greenplum Known Limitations (10.4.0).....	19
PowerExchange for HBase.....	19
PowerExchange for HBase (10.4.1).....	19
PowerExchange for HBase (10.4.0).....	19
PowerExchange for HDFS (10.4.1).....	19
PowerExchange for HDFS Fixed Limitations (10.4.1).....	19
PowerExchange for HDFS Known Limitations (10.4.1).....	20
PowerExchange for HDFS (10.4.0).....	20
PowerExchange for HDFS Fixed Limitations (10.4.0).....	20
PowerExchange for Hive.....	22
PowerExchange for Hive (10.4.1).....	22
PowerExchange for Hive (10.4.0).....	22
PowerExchange for JD Edwards EnterpriseOne.....	23
PowerExchange for JD Edwards EnterpriseOne (10.4.1).....	23
PowerExchange for JD Edwards EnterpriseOne (10.4.0).....	23
PowerExchange for JDBC V2.....	24
PowerExchange for JDBC V2 (10.4.1).....	24
PowerExchange for JDBC V2 Third-Party Known Limitations (10.4.0).....	24
PowerExchange for LDAP.....	24
PowerExchange for LDAP (10.4.1).....	24
PowerExchange for LDAP (10.4.0).....	25
PowerExchange for LinkedIn.....	25
PowerExchange for LinkedIn (10.4.0).....	25
PowerExchange for MapR-DB.....	26
PowerExchange for MapR-DB (10.4.1).....	26
PowerExchange for MapR-DB (10.4.0).....	26
PowerExchange for Microsoft Azure Blob Storage.....	26
PowerExchange for Microsoft Azure Blob Storage (10.4.1).....	26
PowerExchange for Microsoft Azure Blob Storage (10.4.0).....	26
PowerExchange for Microsoft Azure Cosmos DB SQL API.....	27
PowerExchange for Microsoft Azure Cosmos DB SQL API (10.4.1).....	27

PowerExchange for Microsoft Azure Cosmos DB SQL API (10.4.0)	28
PowerExchange for Microsoft Azure Data Lake Storage Gen1	28
PowerExchange for Microsoft Azure Data Lake Storage Gen1 (10.4.1)	28
PowerExchange for Microsoft Azure Data Lake Storage Gen1 (10.4.0)	28
PowerExchange for Microsoft Azure Data Lake Storage Gen2	29
PowerExchange for Microsoft Azure Data Lake Storage Gen2 (10.4.1)	29
PowerExchange for Microsoft Azure Data Lake Storage Gen2 (10.4.0)	31
PowerExchange for Microsoft Azure SQL Data Warehouse	31
PowerExchange for Microsoft Azure SQL Data Warehouse (10.4.1)	31
PowerExchange for Microsoft Azure SQL Data Warehouse (10.4.0)	32
PowerExchange for Microsoft Dynamics CRM	33
PowerExchange for Microsoft Dynamics CRM (10.4.1)	33
PowerExchange for Microsoft Dynamics CRM (10.4.0)	33
PowerExchange for MongoDB	33
PowerExchange for MongoDB (10.4.1)	33
PowerExchange for MongoDB (10.4.0)	33
PowerExchange for Netezza	35
PowerExchange for Netezza (10.4.1)	35
PowerExchange for Netezza (10.4.0)	35
PowerExchange for OData	35
PowerExchange for OData (10.4.1)	35
PowerExchange for OData (10.4.0)	35
PowerExchange for Salesforce	36
PowerExchange for Salesforce (10.4.1)	36
PowerExchange for Salesforce (10.4.0)	36
PowerExchange for Salesforce Marketing Cloud	37
PowerExchange for Salesforce Marketing Cloud (10.4.1)	37
PowerExchange for Salesforce Marketing Cloud (10.4.0)	37
PowerExchange for SAP NetWeaver	37
PowerExchange for SAP NetWeaver (10.4.1)	37
PowerExchange for SAP NetWeaver (10.4.0)	38
PowerExchange for Snowflake	39
PowerExchange for Snowflake (10.4.1)	39
PowerExchange for Snowflake (10.4.0)	40
PowerExchange for Tableau	41
PowerExchange for Tableau (10.4.1)	41

PowerExchange for Tableau (10.4.0)	41
PowerExchange for Tableau V3	42
PowerExchange for Tableau V3 (10.4.1)	42
PowerExchange for Tableau V3 (10.4.0)	42
PowerExchange for Teradata Parallel Transporter API	42
PowerExchange for Teradata Parallel Transporter API (10.4.1)	42
PowerExchange for Teradata Parallel Transporter API (10.4.0)	42
PowerExchange for Twitter	43
PowerExchange for Twitter (10.4.0)	43
PowerExchange for Web Content-Kapow Katalyst	43
PowerExchange for Web Content-Kapow Katalyst (10.4.0)	43
Informatica Global Customer Support	43

Read the *Informatica® PowerExchange Adapters Release Notes* to learn about installation, closed enhancements, fixed limitations, and known limitations for PowerExchange adapters for Informatica.

Note: Effective in version 10.4.1, Cloudera Data Platform as a compute cluster to execute data engineering jobs in the Hadoop environment is available for technical preview. Technical preview functionality is supported for evaluation purposes but is unwarranted and is not production-ready. Informatica recommends that you use in non-production environments only. Informatica intends to include the preview functionality in an upcoming release for production use, but might choose not to in accordance with changing market or technical circumstances. For more information, contact Informatica Global Customer Support.

PowerExchange for Amazon Redshift

PowerExchange for Amazon Redshift (10.4.1)

PowerExchange for Amazon Redshift Fixed Limitations (10.4.1)

There are no fixed limitations in this release.

PowerExchange for Amazon Redshift Known Limitations (10.4.1)

The following table describes known limitations:

Bug	Description
OCON-25666	When you configure an Amazon S3 staging bucket in a different region from the Redshift cluster in a mapping and run the mapping on the Spark engine, the mapping fails with the following error: Mapping execution failed with the following error: File could not be opened because of the following error: [org.apache.hadoop.fs.s3a.AWSS3IOException: getFileStatus on This issue occurs when you use the HDP 2.6.x distribution.

PowerExchange for Amazon Redshift (10.4.0.1)

PowerExchange for Amazon Redshift Fixed Limitations (10.4.0.1)

Review the Release Notes of previous releases for information about previous fixed limitations.

The following table describes fixed limitations:

Bug	Description
OCON-22567	You cannot specify an Amazon Resource Name (ARN) and enable cross-account access when you create a PowerExchange for Amazon Redshift connection.

PowerExchange for Amazon Redshift Known Limitations (10.4.0.1)

There are no known limitations in this release.

PowerExchange for Amazon Redshift (10.4.0)

PowerExchange for Amazon Redshift Fixed Limitations (10.4.0)

Review the Release Notes of previous releases for information about previous fixed limitations.

The following table describes fixed limitations:

Bug	Description
OCON-18377	When you run an Amazon Redshift mapping on the Spark engine, the Data Integration Service executes the post-SQL commands before writing the data to the target.

PowerExchange for Amazon Redshift Known Limitations (10.4.0)

The following table describes known limitations:

Bug	Description
OCON-23228	<p>When you run a mapping that creates a Amazon Redshift target at run time and column names from the output ports of a transformation in the pipeline contain mix of upper case and lower case, the mapping fails with the null pointer exception.</p> <p>Workaround: Use an expression or a case converter transformation to convert the column names into lower case before you select Create Target.</p>
OCON-10209	<p>When you use the MapR distribution, Amazon Redshift mapping fails on the Spark engine when it reads from or writes to an Amazon Redshift cluster that has Version 4 authentication with the following error message:</p> <pre>com.amazonaws.services.s3.model.AmazonS3Exception:</pre>
OCON-9834	<p>When you use the Hortonworks 2.6 distribution, Amazon Redshift mapping fails on the Spark engine when it reads from or writes to an Amazon Redshift cluster that has Version 4 authentication with the following error message:</p> <pre>Bad Request (Service: Amazon S3; Status Code: 400; Error Code: 400 Bad Request; Request ID: 9BDDEEB8241688A2)</pre>
OCON-9827	<p>When you use the IBM BigInsight 4.2 distribution, Amazon Redshift mapping fails on the Spark engine when it reads from or writes to an Amazon Redshift cluster that has Version 4 authentication with the following error message:</p> <pre>com.amazonaws.services.s3.model.AmazonS3Exception: Status Code: 400, AWS Service: Amazon S3, AWS Request ID: 5EEB36DFAC18DE3B, AWS Error Code: null, AWS Error Message: Bad Request</pre>
OCON-9663	<p>When you run an Amazon Redshift mapping to read or write data, the Stop on Errors property does not work.</p>
OCON-8022	<p>If you import an Amazon Redshift table that has a single quote (') in the column name, the mapping fails with the following error message:</p> <pre>[LDTM_0072] [Amazon] (500051) ERROR processing query/statement. Error: Parsing failed, Query: unload ('SELECT "adpqa"."sq_col"."id" FROM "adpqa"."sq_col") TO 's3://infa.qa.bucket/0b0ad503-1c2c-4514-95ac-85a5adb71b3b1489385038407/sq_col_' credentials 'aws_access_key_id=*****;aws_secret_access_key=*****' ESCAPE DELIMITER ','</pre>
OCON-7965	<p>When you run an Amazon Redshift mapping on the Blaze engine to read data from or write data to an Amazon Redshift cluster that requires Version 4 authentication, the mapping fails. This issue occurs if you use the Hortonworks 2.3 distribution.</p>
OCON-7909	<p>When you run an Amazon Redshift mapping on the Blaze engine to read data from or write data to an Amazon Redshift cluster that requires Version 4 authentication, the mapping fails. This issue occurs if you use the MapR 5.2 distribution.</p>
OCON-7322	<p>If you import an Amazon Redshift table that has a single quote (') or a backslash (\) in the table name, the read and write operations fail.</p>

Bug	Description
OCON-6929	<p>If you do not connect all the ports in an Amazon Redshift mapping and run the mapping on the Hive engine, the mapping fails with the following error:</p> <pre>java.lang.RuntimeException [HIVE_1070]The Integration Service failed to run Hive query [exec3_query_2] for task [exec3] due to following error: Hive error code [10,044] , Hive message [FAILED: SemanticException [Error 10044]</pre>
OCON-6921	<p>When you run an Amazon Redshift mapping that contains a timestamp field in the native environment, the Data Integration Service truncates the ultraseconds values to milliseconds.</p>
OCON-6883	<p>When you use IBM BigInsight 4.2 distribution, Amazon Redshift mapping fails on the Blaze engine when it reads from or writes to an Amazon Redshift cluster that has Version 4 authentication.</p>
OCON-6785	<p>When the Amazon Redshift source contains both double quotes (") and the delimiter you specified in the mapping, double quotes are truncated in the target. Also, the escape character is retained in the target.</p>
OCON-6583	<p>If you set the Parallel option off in the unload command and run an Amazon Redshift mapping on the Blaze engine, all the rows from the source are not written to the Amazon Redshift target even though the mapping runs successfully.</p>
OCON-6505	<p>If you specify user impersonation in a Hadoop Connection and run an Amazon Redshift mapping on the Hive engine, no data is written to the target even though the mapping runs successfully.</p>
OCON-6346	<p>When you run an Amazon Redshift mapping on the Blaze engine, the success and error files are not generated.</p>
OCON-6266	<p>When you run an Amazon Redshift mapping that compresses the staging files on the Blaze engine, the mapping fails. The staging files compression is ignored.</p>
OCON-6260	<p>When you run an Amazon Redshift mapping on the Blaze engine, the tasklet log does not display the row statistics even if the mapping runs successfully.</p>
OCON-6252	<p>When you run a mapping on the Blaze engine, the Real and Double data type values are rounded off. Workaround: Use the Numeric data type in place of Real and Double data types.</p>
OCON-1297	<p>When you configure the following attributes and run an Amazon Redshift mapping in the Hadoop environment, the mapping might fail based on the engine selected for mapping execution:</p> <p>Read Operation Attributes:</p> <ul style="list-style-type: none"> - S3 Client Encryption - Staging Directory Location <p>Write Operation Attributes:</p> <ul style="list-style-type: none"> - Enable Compression - CopyOptions Property File - Null value for CHAR and VARCHAR data types - S3 Server Side Encryption - S3 Client Side Encryption - Staging Directory Location - Success File Directory - Error File Directory

Bug	Description
OCON-1275	A mapping with more than one RedShift object fails in Hadoop run-time environment for MapR distribution.
OCON-361	For Hadoop cluster that uses Kerberos, an Amazon Redshift mapping fails on the Hive engine.

PowerExchange for Amazon S3

PowerExchange for Amazon S3 (10.4.1)

PowerExchange for Amazon S3 Fixed Limitations (10.4.1)

The following table describes fixed limitations:

Bug	Description
OCON-25172	When you do a data preview or run a mapping in the dynamic mode for gzip compressed files, you get NULL values.
OCON-24956	When you create a dynamic mapping for flat files and apply a joiner or filter transformation with the condition containing a port name, the mapping fails with the following error: The mapping cannot be validated because of the following errors: [The filter condition in transformation [Filter] is invalid.
OCON-24878	Refresh schema is not applicable to the mappings that run on the Databricks Spark engine.
OCON-24411	When you run a mapping and the source is empty or the source rows are rejected, data is not sent to the target and the Amazon S3 file is empty.
OCON-24130	When you configure the Row delimiter field in an Amazon S3 target file as LF, the target file is generated with the CR LF line endings.

PowerExchange for Amazon S3 Known Limitations (10.4.1)

The following table describes known limitations:

Bug	Description
OCON-25687	When you run a mapping on the Spark engine to read data from a Parquet file, the mapping fails with the following error: Spark task failed with the following error: [com.informatica.sdk.dtm.ExecutionException This issue occurs when you use the CDP 7.1 distribution.
OCON-25683	When you set the <code>-DINFA_HADOOP_DIST_DIR=hadoop\<Distro></code> option in the <code>developerCore.ini</code> file and import a Parquet file, the format of the imported metadata differs based on the distribution. For Cloudera CDP 7.1, the metadata is imported as string and for other supported distributions, the metadata is imported as UTF8.

Bug	Description
OCON-25450	When you create an Amazon S3 data object for the China (Hong Kong) region and run a mapping on the Databricks 5.5 engine, the mapping fails with the following error: Exception caught when redacting in logging: "null". Log original message instead java.lang.NullPointerException
OCON-25032	When you create an Amazon S3 data object for the China (Hong Kong) region and run a mapping on the Spark engine, the mapping fails with the <code>java.lang.reflect.InvocationTargetException</code> exception. This issue occurs when you use the EMR 5.29 distribution
OCON-24683	When you run a mapping in the native environment and read a binary file from a directory, the mapping runs successfully. However, the session log shows the following severe message: SEVERE: jrt write failed.

PowerExchange for Amazon S3 (10.4.0.1)

PowerExchange for Amazon S3 Fixed Limitations (10.4.0.1)

Review the Release Notes of previous releases for information about previous fixed limitations.

The following table describes fixed limitations:

Bug	Description
OCON-22783	You cannot specify an Amazon Resource Name (ARN) and enable cross-account access when you create a PowerExchange for Amazon S3 connection.
OCON-22782	You cannot use FileName port when you write to a binary target.

PowerExchange for Amazon S3 Known Limitations (10.4.0.1)

There are no known limitations in this release.

PowerExchange for Amazon S3 (10.4.0)

PowerExchange for Amazon S3 Fixed Limitations (10.4.0)

There are no fixed limitations in this release.

Review the Release Notes of previous releases for information about previous fixed limitations.

PowerExchange for Amazon S3 Known Limitations (10.4.0)

The following table describes known limitations:

Bug	Description
BDM-11443	When you run a mapping on the Spark engine to read data from a compressed Amazon S3 file and if the file does not have extension for the different compression formats, the mapping fails.
OCON-24554	When you refresh the source or target schema at runtime in a dynamic mapping, the values that you specify for the delimiter, text qualifier, and escape character for a flat file are not honored and the default values are used instead. This might lead to unexpected results in the target.
OCON-23273	When you enable Refresh Schema, the mapping fails to run on the Databricks Spark engine.
OCON-20605	When you run a mapping, in the native environment, to read a flat file that has unicode characters, a space, null values, single quotes, or a value that starts with a dollar sign, the Data Integration Service adds double quotes to the values when writing data to the target.
OCON-17443	When you use the Create Target option to create an Amazon S3 target and select Flat as the Resource Format, fields are not getting propagated to the target. Workaround: Create fields manually in the target file and run the mapping.
OCON-12037	When you import an ORC file that contains Binary data type and other data types, a single binary port is created instead of multiple ports.
OCON-12035	When you run a mapping on the Spark engine to read or write ORC files that contains Binary data type, the task fails with the following error message: <pre>(AmazonS3Read, SparkEngine) = java.lang.RuntimeException</pre>
OCON-12022	When you run a mapping on the Spark engine to read an ORC file that contains Timestamp data type, the mapping runs successfully. However, the Data Integration Service truncates the nanosecond values that are more than six digits.
OCON-11874	When you run a mapping on the Spark engine to read from a multiline JSON file and write the data to the target, the mapping runs successfully. However, the Data Integration Service does not write the data to the target.
OCON-10032	When you use the Hortonworks 2.6 distribution, Amazon S3 mapping fails on the Spark engine when it reads from or writes to an Amazon S3 cluster that has Version 4 authentication with the following error message: <pre>com.amazonaws.services.s3.model.AmazonS3Exception: Status Code: 400, AWS Service: Amazon S3, AWS Request ID: B4AC764FDB8DA642, AWS Error Code: null, AWS Error Message: Bad Request</pre>
OCON-10028	When you use the IBM BigInsight 4.2 distribution, Amazon S3 mapping fails on the Spark engine when it reads from or writes to an Amazon S3 cluster that has Version 4 authentication with the following error message: <pre><CmdExecInProgressTasks-pool-2-thread-29> SEVERE:[Pre_Spark_Task_Command 1] [com.informatica.platform.dtm.executor.hadoop.impl.cmdtasks.HadoopFSRmRfTask]</pre>

Bug	Description
OCON-10026	<p>When you use the MapR 5.2 secure cluster, Amazon S3 mapping fails on the Spark engine when it reads from or writes to an Amazon S3 cluster that has Version 4 authentication with the following error message:</p> <pre data-bbox="344 422 1476 478"><CmdExecInProgressTasks-pool-2-thread-29> SEVERE:[Pre_Spark_Task_Command_1] [com.informatica.platform.dtm.executor.hadoop.impl.cmdtasks.HadoopFSRmRfTask]</pre>
OCON-9765	<p>When you read files within a sub-directory that contains different compression formats, the Data Integration Service does not write the data to the target properly.</p>
OCON-9353	<p>When you create an Amazon S3 data object for the US East (Ohio) region and run a mapping on the Spark engine, the task fails. However, the same mapping runs successfully in the native environment.</p>
OCON-8996	<p>Unable to view the list of files available in a bucket when you expand the bucket name list in the Object Explorer view. Workaround: To view the list of files available in a bucket, you must select the bucket name instead of expanding the bucket name list in the Object Explorer view.</p>
OCON-7963	<p>When you run an Amazon S3 mapping on the Blaze engine to read data from or write data to an Amazon S3 bucket that requires Version 4 authentication, the mapping fails. This issue occurs if you use the Hortonworks 2.3 distribution.</p>
OCON-7938	<p>When you run an Amazon S3 mapping in the native environment, the incorrect values in the decimal or bigint data type field are replaced by zeroes. However, the mapping runs successfully and the session log displays the following error message:</p> <pre data-bbox="344 1066 1476 1150">2017-03-13 23:46:08.131 <TASK_140116755179264-READER_1_1_1> SEVERE: [APPSDK_Msg_1762] Data for column [age] of type [bigint] should be a of type [java.lang.Number] or its sub-types.</pre>
OCON-7911	<p>When you run an Amazon S3 mapping on the Blaze engine to read data from or write data to an Amazon S3 bucket that requires Version 4 authentication, the mapping fails. This issue occurs if you use the MapR 5.2 distribution.</p>
OCON-7743	<p>When you connect to Amazon S3 from the Administrator console for a Kerberos domain, database TLS enabled, or enabled secure communication, the test connection fails with the following error message:</p> <pre data-bbox="344 1396 1476 1453">The requested operation could not be performed due to the following error: Connection error: Unable to execute HTTP reuest:</pre>
OCON-7024	<p>An Amazon S3 mapping fails on the Blaze engine when it reads from or write to an Amazon S3 cluster that uses IBM BigInsight 4.2 distribution and requires Version 4 authentication.</p>
OCON-6887	<p>If the delimiter in the source and the target is a multibyte character, the PowerExchange for Amazon S3 mapping fails.</p>
OCON-5736	<p>When you run an Amazon S3 mapping that writes to an Amazon S3 target on the Blaze engine, the hash symbol (#) precedes the header in the Amazon S3 target.</p>

PowerExchange for Cassandra

PowerExchange for Cassandra (10.4.1)

There are no fixed or known limitations in this release.

Review the Release Notes of previous releases for information about previous fixed limitations.

PowerExchange for Cassandra (10.4.0)

PowerExchange for Cassandra Fixed Limitations (10.4.0)

There are no fixed limitations in this release.

Review the Release Notes of previous releases for information about previous fixed limitations.

PowerExchange for Cassandra Known Limitations (10.4.0)

The following table describes known limitations:

Bug	Description
OCON-11018	When you select the Show Default Schema option in the Connection Explorer view, the Informatica Cassandra ODBC driver fails to fetch the tables from the default keyspace. Workaround: Disable the Show Default Schema option in the connection explorer and browse the keyspace to import the metadata. (421174)

PowerExchange for Cassandra Third-Party Known Limitations (10.4.0)

The following table describes third-party known limitations:

Bug	Description
OCON-11307	The Cassandra database stores values of the Float data type in the exponential format. The value of the Float data type read by the Informatica Cassandra ODBC driver is not the same as the value of the Float data type in the Cassandra database because the Informatica Cassandra ODBC driver maps the Float data type to SQL_Real. (397018)
OCON-11206	The Cassandra database does not allow periods in column names. PreSQL or PostSQL queries generated with the SQL Editor fail with a syntax error because the SQL Editor prefixes periods to column names to qualify column names with table names. Workaround: Do not qualify column names with table names in queries. (397187)
OCON-7787	The performance of reading data from a Cassandra source is slow even when you increase the number of partitions.
OCON-70237	The Cassandra ODBC driver does not support null values in collections and reports non-key columns in virtual tables as not nullable.
OCON-70231	If you use a single SQL statement to write multiple rows to a Cassandra target and you bind the default value to the first row, the insert statement fails. Workaround: Use null instead of the default value.

PowerExchange for Cassandra JDBC

PowerExchange for Cassandra JDBC (10.4.1)

There are no fixed or known limitations in this release.

Review the Release Notes of previous releases for information about previous fixed limitations.

PowerExchange for Cassandra JDBC (10.4.0)

PowerExchange for Cassandra JDBC Fixed Limitations (10.4.0)

There are no fixed limitations in this release.

Review the Release Notes of previous releases for information about previous fixed limitations.

PowerExchange for Cassandra JDBC Known Limitations (10.4.0)

The following table describes known limitations:

Bug	Description
OCON-11018	When you select the Show Default Schema option in the Connection Explorer view, the Informatica Cassandra ODBC driver fails to fetch the tables from the default keyspace. Workaround: Disable the Show Default Schema option in the connection explorer and browse the keyspace to import the metadata. (421174)

PowerExchange for Cassandra JDBC Third-Party Fixed Limitations (10.4.0)

The following table describes third-party known limitations:

Bug	Description
OCON-18017	When you run a mapping to read data of INET data type in a Cassandra source table and if the column does not contain any data, the mapping fails. Simba ticket reference number: 00133187

PowerExchange for DataSift

PowerExchange for DataSift (10.4.0)

There are no fixed or known limitations in this release.

Review the Release Notes of previous releases for information about previous fixed limitations.

PowerExchange for Facebook

PowerExchange for Facebook (10.4.0)

PowerExchange for Facebook Fixed Limitations (10.4.0)

There are no fixed limitations in this release.

Review the Release Notes of previous releases for information about previous fixed limitations.

PowerExchange for Facebook Known Limitations (10.4.0)

The following table describes known limitations:

Bug	Description
OCON-9933	<p>The Scope attribute in PowerExchange for Facebook connection properties shows incorrect values.</p> <p>Workaround: Remove the following values from the Scope attribute before you test the connection or get access tokens:</p> <ul style="list-style-type: none">- user_activities- user_groups- user_interests- manage_notifications- read_friendlists- read_mailbox- read_stream <p>For more details about Facebook login permissions, see the Facebook documentation.</p>

PowerExchange for Google Analytics

PowerExchange for Google Cloud Analytics (10.4.1)

There are no fixed or known limitations in this release.

Review the Release Notes of previous releases for information about previous fixed limitations.

PowerExchange for Google Analytics (10.4.0)

PowerExchange for Google Analytics Fixed Limitations (10.4.0)

There are no fixed limitations in this release.

Review the Release Notes of previous releases for information about previous fixed limitations.

PowerExchange for Google Analytics Known Limitations (10.4.0)

The following table describes known limitations:

Bug	Description
OCON-18405	When you use a PowerExchange for Google Analytics data object read operation to read data from a Google Analytics report and select a field that contains <code>xx</code> in the field name, the Data Integration Service fails to fetch the metadata from the field and the mapping task fails with the following error: <pre>[ERROR] java.lang.NullPointerException</pre>

PowerExchange for Google BigQuery

PowerExchange for Google BigQuery (10.4.1)

PowerExchange for Google BigQuery Fixed Limitations (10.4.1)

There are no fixed limitations in this release.

Review the Release Notes of previous releases for information about previous fixed limitations.

PowerExchange for Google BigQuery Known Limitations (10.4.1)

The following table describes known limitations:

Bug	Description
OCON-25438	When you configure an uncached lookup operation to look up data from a Google BigQuery source, the session fails when one of the following conditions is true: <ul style="list-style-type: none">- You select Direct for the Read Mode property.- You set the Number of Rows to Read property to a non-zero value.- You set the On multiple matches property to Return Any.- You configure the SQL Override Query property.

PowerExchange for Google BigQuery Third-Party Known Limitations (10.4.1)

The following table describes third-party known limitations:

Bug	Description
OCON-25412	<p>When you run a Google BigQuery mapping on the Spark engine, the mapping fails when the following conditions are true:</p> <ul style="list-style-type: none">- You configure a Google BigQuery source and select Optimized as the Spark Mode.- You configure a Google BigQuery target and select Optimized as the Spark Mode.- You use any one of the following versions of the Hadoop distribution in the Hadoop environment:<ul style="list-style-type: none">- Amazon EMR 5.29- Amazon EMR 5.26- Azure HDInsight 3.6- Cloudera CDH 5.16- Cloudera CDH 5.13- Hortonworks HDP 2.6- MapR <p>Google ticket reference number: 157261131</p>

PowerExchange for Google BigQuery (10.4.0)

PowerExchange for Google BigQuery Fixed Limitations (10.4.0)

There are no fixed limitations in this release.

Review the Release Notes of previous releases for information about previous fixed limitations.

PowerExchange for Google BigQuery Known Limitations (10.4.0)

The following table describes known limitations:

Bug	Description
OCON-22940	If you use the Create Target option to create a Google BigQuery target and the source contains a column of Integer data type, the Data Integration Service fails to create a Google BigQuery target.
OCON-17956	The test connection is successful even when you specify incorrect credentials in the connection properties.

PowerExchange for Google BigQuery Third-Party Known Limitations (10.4.0)

The following table describes third-party known limitations:

Bug	Description
OCON-22676	<p>When you run a mapping to read data of timestamp data type from a Google BigQuery source, incorrect values are written to the target for certain timestamp values.</p> <p>Google ticket reference number: 142002729</p>

PowerExchange for Google Cloud Spanner

PowerExchange for Google Cloud Spanner (10.4.1)

There are no fixed or known limitations in this release.

Review the Release Notes of previous releases for information about previous fixed limitations.

PowerExchange for Google Cloud Spanner (10.4.0)

There are no fixed or known limitations in this release.

Review the Release Notes of previous releases for information about previous fixed limitations.

PowerExchange for Google Cloud Storage

PowerExchange for Google Cloud Storage (10.4.1)

PowerExchange for Google Cloud Storage Fixed Limitations (10.4.1)

Review the Release Notes of previous releases for information about previous fixed limitations.

The following table describes fixed limitations:

Bug	Description
OCON-23109	If you select the Is Directory option in the data object read operation advanced properties to read files in a Google Cloud Storage folder in JSON or Binary format and run a mapping in the native environment, the mapping fails.

PowerExchange for Google Cloud Storage Known Limitations (10.4.1)

The following table describes known limitations:

Bug	Description
OCON-25683	When you set the <code>-DINFA_HADOOP_DIST_DIR=hadoop\<Distro></code> option in the <code>developerCore.ini</code> file and import a Parquet file, the format of the imported metadata differs based on the distribution. For Cloudera CDP 7.1, the metadata is imported as string and for other supported distributions, the metadata is imported as UTF8.

PowerExchange for Google Cloud Storage (10.4.0)

PowerExchange for Google Cloud Storage Fixed Limitations (10.4.0)

Review the Release Notes of previous releases for information about previous fixed limitations.

The following table describes fixed limitations:

Bug	Description
OCON-18427	<p>When you run a mapping on the Spark Engine to write data to a Google Cloud Storage flat file target and the data size is greater than 750 MB, the mapping hangs indefinitely with the following error:</p> <pre>ERROR AbstractGoogleAsyncWriteChannel: Exception not convertible into handled response java.io.IOException: Write end dead</pre> <p>Workaround: Set the following property in the Spark Execution Parameters in the Hadoop connection configured for the Spark engine:</p> <pre>spark.executor.instances=1</pre> <p>Note: When you run the mapping on a single spark executor instance, you may experience performance degradation.</p>
OCON-17842	<p>When you use the Create Target option to create a Google Cloud Storage target and select Flat as the Resource Format, fields are not getting propagated to the target.</p> <p>Workaround: Create fields manually in the target file and run the mapping.</p>

PowerExchange for Google Cloud Storage Known Limitations (10.4.0)

The following table describes known limitations:

Bug	Description
OCON-23109	<p>If you select the Is Directory option in the data object read operation advanced properties to read files in a Google Cloud Storage folder in JSON or Binary format and run a mapping in the native environment, the mapping fails.</p>
OCON-17932	<p>When you specify the Google Cloud Storage Path in the <code>gs://<bucket name></code> format and run a mapping on the Spark engine to write data to a Google Cloud Storage target, the mapping fails.</p> <p>Workaround: Specify the Google Cloud Storage Path in the following format:</p> <pre>gs://<bucket name>/<folder_name></pre>

PowerExchange for Greenplum

PowerExchange for Greenplum (10.4.1)

There are no fixed or known limitations in this release.

Review the Release Notes of previous releases for information about previous fixed limitations.

PowerExchange for Greenplum (10.4.0)

PowerExchange for Greenplum Fixed Limitations (10.4.0)

There are no fixed limitations in this release.

Review the Release Notes of previous releases for information about previous fixed limitations.

PowerExchange for Greenplum Known Limitations (10.4.0)

The following table describes known limitations:

Bug	Description
PLAT-13542	When you create a Greenplum connection by using the <code>infacmd isp CreateConnection</code> command, you must enter even the optional fields.

PowerExchange for HBase

PowerExchange for HBase (10.4.1)

There are no fixed or known limitations in this release.

Review the Release Notes of previous releases for information about previous fixed limitations.

PowerExchange for HBase (10.4.0)

PowerExchange for HBase Fixed Limitations (10.4.0)

Review the Release Notes of previous releases for information about previous fixed limitations.

The following table describes fixed limitations:

Bug	Description
OCON-22607	When you run an HBase mapping in the native environment and on the Spark engine to read <code>DateTime</code> field data, the Data Integration Service returns different time zone values from a given timestamp in the native and Hadoop environments.
OCON-21401	When you run a mapping to read data from an HBase source table and the metadata contains a space in column family names, the mapping fails.
OCON-21400	When you run a mapping to read data from an HBase source table and the metadata contains columns in binary format, the mapping fails.

PowerExchange for HBase Known Limitations (10.4.0)

There are no known limitations for this release.

PowerExchange for HDFS (10.4.1)

PowerExchange for HDFS Fixed Limitations (10.4.1)

There are no fixed limitations in this release.

Review the Release Notes of previous releases for information about previous fixed limitations.

PowerExchange for HDFS Known Limitations (10.4.1)

The following table describes known limitations:

Bug	Description
OCON-25683	When you set the <code>-DINFA_HADOOP_DIST_DIR=hadoop\<Distro></code> option in the <code>developerCore.ini</code> file and import a Parquet file, the format of the imported metadata differs based on the distribution. For Cloudera CDP 7.1, the metadata is imported as string and for other supported distributions, the metadata is imported as UTF8.
OCON-25655	When you run a mapping in the native environment to write data to a complex file target and choose to overwrite the target data and the target filename does not contain the file format extension such as ".avro" or ".parquet", the Data Integration Service does not delete the target data before writing data.
OCON-25649	When you run a mapping on the Blaze engine to read data from and write data to sequence complex file data objects using a Kerberos connection, the mapping fails with the following exception: "GSSException: No valid credentials provided (Mechanism level: Failed to find any Kerberos target)"
OCON-24348	When you run a mapping on the Spark engine to read data from an empty JSON complex file source and write data to a complex file target, the mapping should fail, but the mapping runs successfully and the Data Integration Service generates an empty target file.

PowerExchange for HDFS (10.4.0)

PowerExchange for HDFS Fixed Limitations (10.4.0)

Review the Release Notes of previous releases for information about previous fixed limitations.

The following table describes fixed limitations:

Bug	Description
OCON-19803	When you run a mapping on the Spark engine to read data from a local complex file source, if the folders inside the source directory contains files with same names and you use a wildcard pattern to read the entire parent directory, the Data Integration Service reads data from any one of the files with the same names and there is no impact in reading from the other files.
OCON-18292	When you use the <code>FileName</code> port in a complex file target in Binary format and run a mapping in the native environment, the Data Integration Service writes all the files to a single folder.
OCON-18208	When you run a mapping to write data to a complex file using Create target option for Avro/Parquet formats with mapping flow enabled, the schema is created with primitive datatypes and the Data Integration Service skips the rows having null values.

Bug	Description
OCON-18169	When you run a mapping on the Spark engine to read data from a complex file with ORC (Optimized Row Columnar) datatype using the wildcard character questionmark '?' in MapR distro, the mapping fails.
BDM-14811	When you validate a mapping and select a connection parameter type that is not valid, the parameter name appears incorrectly. This error occurs when you import a flat file from the Hadoop environment, parameterize the connection name, and change the parameter type to a type that is not valid.

PowerExchange for HDFS Known Limitations (10.4.0)

The following table describes known limitations:

Bug	Description
OCON-23230	You cannot use multiple level partitioning when you run a mapping to write data to a complex file target with Filename port enabled.
OCON-23124	When you run a mapping in the native environment to write data to a complex file using filename port and mapping flow is enabled, the Data Integration Service generates an incorrect folder structure and writes data to a single file.
OCON-23122	When you run a mapping in the native environment to read or write data to a complex file object in ORC format, the mapping fails
OCON-23084	When you run a mapping to read from a complex file source and write to a complex file target and the source object contains unsupported data types in the schema, the mapping fails.
OCON-21852	When you import a complex file data object in JSON format, the import fails with the following error: <code>Array must contain at least 1 element for projection</code>
OCON-17103	When you run a mapping on the Spark engine to read data from a complex file and the source path has a wildcard character, then the log file does not display the source file names.
OCON-16280	When you create a complex file data object from a JSON file, the task fails with the following error: <code>Encountered an error saving the data object</code>
OCON-15862	When you run a mapping to read data from a complex file source in JSON format and enable compression, the mapping runs successfully but the Data Integration Service fails to read data from source.
OCON-12579	If you set the Hive warehouse directory in a Hadoop connection to an encrypted HDFS directory and the impersonation user does not have the DECRYPT_EEK permission, complex file mappings run indefinitely on the Hive engine.

PowerExchange for Hive

PowerExchange for Hive (10.4.1)

PowerExchange for Hive Fixed Limitations (10.4.1)

Review the Release Notes of previous releases for information about previous fixed limitations.

The following table describes fixed limitations:

Bug	Description
OCON-24896	When you run a mapping in the native environment to write data to a Hive target, the mapping runs successfully. However, the session log displays the following error: <code>java.lang.ArrayIndexOutOfBoundsException</code>

PowerExchange for Hive Known Limitations (10.4.1)

The following table describes known limitations:

Bug	Description
BDM-33998	When you run a mapping on the Spark engine to write data to a Hive table with complex data types such as struct and the column name contains a reserved word, the mapping fails.
BDM-33990	When you run a mapping on the Spark engine to write data to a Hive table with complex data types such as struct and the column names have special characters, the mapping fails.
OCON-25634	When you alter a Hive table by adding a new column to a Hive partition table, the mapping fails with the following validation error: <code>"You cannot add column {columnname} as table {tablename} is partitioned"</code>
OCON-25343	When you change the table metadata in the Developer Tool by modifying the existing data type in a Hive mapping that reads data from a Hive source and do not synchronize the Physical Data Object (PDO), the mapping fails with the following error: <code>FAILED: SemanticException [Error 10044]: Line 1:23 Cannot insert into target table because column number/types are different</code> Workaround: Synchronize the Physical Data Object and run the mapping again.
OCON-25211	When you run a mapping to read data from a Hive table with hierarchical (Htype) data type in columns and use the sort option to override the default SQL query, the mapping fails.
OCON-25180	When you synchronize a Hive object that contains complex datatypes in the Developer tool, the links between the hive objects in the mapping are not retained.

PowerExchange for Hive (10.4.0)

PowerExchange for Hive Fixed Limitations (10.4.0)

Review the Release Notes of previous releases for information about previous fixed limitations.

The following table describes fixed limitations:

Bug	Description
OCON-16162	When you import a Hive table and run a mapping the custom query button is unable to render object and gives following error: File (ValidateButtonQueryHiveNotFunctional.JPG)not found.File ValidateButtonDropDownNothingsFunctional.JPG)not found.

PowerExchange for Hive Known Limitations (10.4.0)

The following table describes known limitations:

Bug	Description
OCON-18287	When you alter a Hive table by adding a new column in native environment and the table contains a special character in the column name, the mapping fails.
OCON-17109	When you run a Hive mapping in the native environment with decimal data type in the source table having precision more than 28 and high precision mode is enabled, the mapping runs successfully but the data is not written correctly in target.
BDM-16907	When you specify a custom JDBC driver class name, the Hive mapping fails on the Blaze and the Spark engines with an <code>UndeclaredThrowableException</code> error.

PowerExchange for JD Edwards EnterpriseOne

PowerExchange for JD Edwards EnterpriseOne (10.4.1)

There are no fixed or known limitations in this release.

Review the Release Notes of previous releases for information about previous fixed limitations.

PowerExchange for JD Edwards EnterpriseOne (10.4.0)

PowerExchange for JD Edwards EnterpriseOne Fixed Limitations (10.4.0)

There are no fixed limitations in this release.

Review the Release Notes of previous releases for information about previous fixed limitations.

PowerExchange for JD Edwards EnterpriseOne Known Limitations (10.4.0)

The following table describes known limitations:

Bug	Description
OCON-11523	When you import a table that contains the ID_LONG data type, the Integration Service fails to import some of the table columns.
OCON-11456	Even when the error threshold is reached, the Integration Service continues to process the data and mapping does not fail. The issue occurs because the Stop On Errors run-time property does not work.
OCON-11372	If the Integration Service writes rejected rows followed by valid rows to an interface table, the row statistics generated for an InterfaceWrite operation is incorrect.
OCON-11223	When you apply a native filter expression for data that contains the JDE date data type format that is not valid, the data preview and mapping does not fail. An appropriate error message does not appear when you use formats that are not valid. Workaround: Do not use the yyyy-mm-dd hh:mm:ss and hh:mm:ss date formats. Instead, use the yyyy-mm-dd format.

PowerExchange for JDBC V2

PowerExchange for JDBC V2 (10.4.1)

There are no fixed or known limitations in this release.

Review the Release Notes of previous releases for information about previous fixed limitations.

PowerExchange for JDBC V2 Third-Party Known Limitations (10.4.0)

The following table describes known limitations:

Bug	Description
OCON-22649	When you run a JDBC V2 mapping to read or write Time data type with precision greater than 3, data is truncated.

PowerExchange for LDAP

PowerExchange for LDAP (10.4.1)

There are no fixed or known limitations in this release.

Review the Release Notes of previous releases for information about previous fixed limitations.

PowerExchange for LDAP (10.4.0)

PowerExchange for LDAP Fixed Limitations (10.4.0)

There are no fixed limitations in this release.

Review the Release Notes of previous releases for information about previous fixed limitations.

PowerExchange for LDAP Known Limitations (10.4.0)

The following table describes known limitations:

Bug	Description
OCON-11513	When you run a mapping, the Integration Service communicates with the LDAP server directly instead of making or receiving calls through the proxy server.
OCON-11338	When you enable CDC, you cannot fetch entries related to deleted records. The error occurs when the recycle bin on the Active Directory server is full and cannot store more deleted records. Workaround: Empty the recycle bin on the Active Directory server and try again.
OCON-6882	On Solaris, a mapping fails when you perform a lookup against an LDAP source.
OCON-1151	Even when the error threshold is reached, the Integration Service continues to process the data and the mapping does not fail. The issue occurs because the Stop On Errors run-time property does not work.

PowerExchange for LinkedIn

PowerExchange for LinkedIn (10.4.0)

PowerExchange for LinkedIn Fixed Limitations (10.4.0)

There are no fixed limitations in this release.

Review the Release Notes of previous releases for information about previous fixed limitations.

PowerExchange for LinkedIn Known Limitations (10.4.0)

The following table describes known limitations:

Bug	Description
OCON-9941	When you create a PowerExchange for LinkedIn connection, getting the OAuth token and secret for LinkedIn might fail. Workaround: When you create a connection, verify that the Scope attribute contains the values defined in the LinkedIn application. The default values specified in the Scope attribute might not work. For more details, see the LinkedIn documentation.

PowerExchange for MapR-DB

PowerExchange for MapR-DB (10.4.1)

There are no fixed or known limitations in this release.

Review the Release Notes of previous releases for information about previous fixed limitations.

PowerExchange for MapR-DB (10.4.0)

There are no fixed or known limitations in this release.

Review the Release Notes of previous releases for information about previous fixed limitations.

PowerExchange for Microsoft Azure Blob Storage

PowerExchange for Microsoft Azure Blob Storage (10.4.1)

PowerExchange for Microsoft Azure Blob Storage Fixed Limitations (10.4.1)

There are no fixed limitations in this release.

PowerExchange for Microsoft Azure Blob Storage Known Limitations (10.4.1)

The following table describes known limitations:

Bug	Description
OCON-25683	When you set the <code>-DINFA_HADOOP_DIST_DIR=hadoop\<Distro></code> option in the <code>developerCore.ini</code> file and import a Parquet file, the format of the imported metadata differs based on the distribution. For Cloudera CDP 7.1, the metadata is imported as string and for other supported distributions, the metadata is imported as UTF8.

PowerExchange for Microsoft Azure Blob Storage (10.4.0)

PowerExchange for Microsoft Azure Blob Storage Fixed Limitations (10.4.0)

There are no fixed limitations in this release.

Review the Release Notes of previous releases for information about previous fixed limitations.

PowerExchange for Microsoft Azure Blob Storage Known Limitations (10.4.0)

The following table describes known limitations:

Bug	Description
BDM-19847	For the write operation, when you run a mapping on the Spark engine and the folder path contains special characters, the Data Integration Service creates a new folder.
OCON-24942	When you refresh the source or target schema at runtime in a dynamic mapping, the values that you specify for the delimiter, text qualifier, and escape character for a flat file are not used and the default values are used. This issue leads to unexpected results in the target.
OCON-22511	When you read data from a Microsoft Azure SQL Data Warehouse source and use the Create Target option to create a Microsoft Azure Blob Storage target, if the Microsoft Azure Blob Storage connection uses SAS authentication, the mapping fails.
OCON-20605	When you run a mapping, in the native environment, to read a flat file that has unicode characters, a space, null values, single quotes, or a value that starts with a dollar sign, the Data Integration Service adds double quotes to the values when writing data to the target.
OCON-17642	<p>When you enable Mapping Flow in a mapping that reads data from a flat file source and writes to a flat file target, the mapping fails with the following error in the native environment:</p> <pre>java.lang.Exception: [MPSVCCMN_10094] The Mapping Service Module failed to run the job with ID [Ic2j9ASPEemTlSYmtVHPww] because of the following error: [EdtmExec_00007] Exception: /tmp/insertd29a7def_bb59_452d_8051_ea4b4630807b9132318161205585091.azb (No such file or directory)</pre> <p>Workaround: Remove the FileName field from the imported source object and rerun the mapping.</p>
OCON-17443	<p>When you use the Create Target option to create a Microsoft Azure Blob Storage target and select Flat as the Resource Format, fields are not getting propagated to the target.</p> <p>Workaround: Enable column projection and create fields manually in the target file and run the mapping.</p>
OCON-17082	When you import an object from sub directories with names having a space, data preview fails.
OCON-12420	When you read or write a blob that has special characters, the mapping fails on the Spark engine.
OCON-12352	When a JSON file contains special characters, the Data Integration Service does not read the data correctly in the Spark mode.
OCON-12318	The Data Integration Service adds an extra blank new line at the end when you read or write a flat file in the native environment or in the Spark mode.
OCON-10125	When you read data from or write data to Microsoft Azure Blob Storage, the entire blob gets downloaded in the staging directory even if you cancel the mapping.

PowerExchange for Microsoft Azure Cosmos DB SQL API

PowerExchange for Microsoft Azure Cosmos DB SQL API (10.4.1)

There are no fixed or known limitations in this release.

Review the Release Notes of previous releases for information about previous fixed limitations.

PowerExchange for Microsoft Azure Cosmos DB SQL API (10.4.0)

PowerExchange for Microsoft Azure Cosmos DB SQL API Fixed Limitations (10.4.0)

There are no fixed limitations in this release.

Review the Release Notes of previous releases for information about previous fixed limitations.

PowerExchange for Microsoft Azure Cosmos DB SQL API Known Limitations (10.4.0)

The following table describes known limitations:

Bug	Description
OCON-11892	When you create a data object, the system fields, for example <code>_ts</code> , <code>_etag</code> , appear by default in the write operation. Workaround: Go to Column Projection and edit the schema to remove the system fields.

PowerExchange for Microsoft Azure Data Lake Storage Gen1

PowerExchange for Microsoft Azure Data Lake Storage Gen1 (10.4.1)

PowerExchange for Microsoft Azure Data Lake Storage Gen1 Fixed Limitations (10.4.1)

There are no fixed limitations in this release.

PowerExchange for Microsoft Azure Data Lake Storage Gen1 Known Limitations (10.4.1)

The following table describes known limitations:

Bug	Description
OCON-25683	When you set the <code>-DINFA_HADOOP_DIST_DIR=hadoop\<Distro></code> option in the <code>developerCore.ini</code> file and import a Parquet file, the format of the imported metadata differs based on the distribution. For Cloudera CDP 7.1, the metadata is imported as string and for other supported distributions, the metadata is imported as UTF8.

PowerExchange for Microsoft Azure Data Lake Storage Gen1 (10.4.0)

PowerExchange for Microsoft Azure Data Lake Storage Gen1 Fixed Limitations (10.4.0)

There are no fixed limitations in this release.

Review the Release Notes of previous releases for information about previous fixed limitations.

PowerExchange for Microsoft Azure Data Lake Storage Gen1 Known Limitations (10.4.0)

The following table describes known limitations:

Bug	Description
OCON-24942	When you refresh the source or target schema at runtime in a dynamic mapping, the values that you specify for the delimiter, text qualifier, and escape character for a flat file are not used and the default values are used. This issue leads to unexpected results in the target.
OCON-20605	When you run a mapping, in the native environment, to read a flat file that has unicode characters, a space, null values, single quotes, or a value that starts with a dollar sign, the Data Integration Service adds double quotes to the values when writing data to the target.
OCON-17443	When you use the Create Target option to create a Microsoft Azure Blob Storage target and select Flat as the Resource Format, fields are not getting propagated to the target. Workaround: Enable column projection and create fields manually in the target file and run the mapping.
OCON-10244	When you read data from or write data to Microsoft Azure Data Lake Store, the entire file gets downloaded in the staging directory even if you cancel the mapping.

PowerExchange for Microsoft Azure Data Lake Storage Gen2

PowerExchange for Microsoft Azure Data Lake Storage Gen2 (10.4.1)

PowerExchange for Microsoft Azure Data Lake Storage Gen2 Fixed Limitations (10.4.1)

Review the Release Notes of previous releases for information about previous fixed limitations.

The following table describes fixed limitations:

Bug	Description
OCON-25503	The test connection fails due to a file system name error.

PowerExchange for Microsoft Azure Data Lake Storage Gen2 Known Limitations (10.4.1)

The following table describes known limitations:

Bug	Description
OCON-25687	<p>When you run a mapping to read data from a Parquet file on the Spark engine, the mapping fails with the following error:</p> <pre>Spark task failed with the following error: [com.informatica.sdk.dtm.ExecutionException</pre> <p>This issue occurs when you use the Cloudera CDP 7.1 distribution.</p>
OCON-25683	<p>When you set the <code>-DINFA_HADOOP_DIST_DIR=hadoop\<Distro></code> option in the <code>developerCore.ini</code> file and import a Parquet file, the format of the imported metadata differs based on the distribution. For Cloudera CDP 7.1, the metadata is imported as string and for other supported distributions, the metadata is imported as UTF8.</p>
OCON-24963	<p>When you run a mapping to write a complex file to a Microsoft Azure Data Lake Storage Gen2 target on Databricks Spark engine, the Data Integration Service does not override all the target files of the previous mapping in the target directory.</p> <p>This issue occurs for all target files except the partition file.</p>
OCON-24923	<p>When you read an ORC file on Spark engine or Databricks Spark engine, the Data Integration Service fails to read the special characters in the ORC file.</p>
OCON-24734	<p>When you read a JSON file in the native environment, the Data Integration Service incorrectly writes the values of double data type to the target in exponential format.</p>
OCON-23135	<p>When you import a flat file data object with default precision and update the precision value at run time, the Data Integration Service uses the default value instead of the value provided at run time.</p>

PowerExchange for Microsoft Azure Data Lake Storage Gen2 Third-Party Known Limitations (10.4.1)

The following table describes third-party known limitations:

Bug	Description
OCON-18625	You cannot configure the Azure Government end-points in mappings that run on Databricks Spark engine.

PowerExchange for Microsoft Azure Data Lake Storage Gen2 (10.4.0)

PowerExchange for Microsoft Azure Data Lake Storage Gen2 Known Limitations (10.4.0)

The following table describes known limitations:

Bug	Description
BDM-28877	When you select <code>No quotes</code> in Text qualifier and specify the escape character in the flat file schema format properties, extra columns are added while previewing the data. Workaround: Select <code>Single quotes</code> or <code>Double quotes</code> in Text qualifier to import the object. After the object import, you can select <code>No quotes</code> in the Schema tab.
OCON-24982	When you use the <code>FileName</code> port in a complex file target and run a mapping with mapping flow enabled in the native environment, the Data Integration Service does not create the folder structure as expected.
OCON-24942	When you refresh the source or target schema at runtime in a dynamic mapping, the values that you specify for the delimiter, text qualifier, and escape character for a flat file are not used and the default values are used. This issue leads to unexpected results in the target.
OCON-23033	When you import an object from sub directories with names having a space, data preview fails.
OCON-21916	When you create a pass-through mapping to read data from a flat file and write data to another flat file target and enable run-time linking, <code>'#'</code> is added to the first column name in the target file when you run the mapping first time on the Spark engine. The mapping fails at subsequent runs.
OCON-20605	When you run a mapping, in the native environment, to read a flat file that has unicode characters, a space, null values, single quotes, or a value that starts with a dollar sign, the Data Integration Service adds double quotes to the values when writing data to the target.

PowerExchange for Microsoft Azure SQL Data Warehouse

PowerExchange for Microsoft Azure SQL Data Warehouse (10.4.1)

PowerExchange for Microsoft Azure SQL Data Warehouse Fixed Limitations (10.4.1)

Review the Release Notes of previous releases for information about previous fixed limitations.

The following table describes fixed limitations:

Bug	Description
OCON-24469	When you read data from a Microsoft Azure SQL Data Warehouse source, the Data Integration Service does not display all the rows in data preview.
OCON-21693	If the Informatica domain is SSL-enabled, dynamic mappings fail in the native environment. Microsoft ticket reference number: 119081423001742

PowerExchange for Microsoft Azure SQL Data Warehouse Known Limitations (10.4.1)

There are no known limitations in this release.

PowerExchange for Microsoft Azure SQL Data Warehouse (10.4.0)

PowerExchange for Microsoft Azure SQL Data Warehouse Fixed Limitations (10.4.0)

Review the Release Notes of previous releases for information about previous fixed limitations.

The following table describes fixed limitations:

Bug	Description
OCON-19620	When you use an ODBC data source to connect to Microsoft Azure SQL Data Warehouse and if the imported metadata contains unicode characters, the mapping fails.

PowerExchange for Microsoft Azure SQL Data Warehouse Known Limitations (10.4.0)

The following table describes known limitations:

Bug	Description
BDM-28917	When the source data contains a slash '\ ' and you create a mapping to read the source data and retain the default quote character, 0x1f, the mapping fails with the following error on the Databricks Spark engine: SEVERE: Exception: com.microsoft.sqlserver.jdbc.SQLServerException: HdfsBridge::recordReaderFillBuffer - Unexpected error encountered filling record reader buffer: HadoopExecutionException: Could not find a delimiter after string delimiter."
OCON-22851	When you use an ODBC data source to connect to Microsoft Azure SQL Data Warehouse and if the imported metadata contains special characters, the mapping fails.
OCON-22639	When you run a dynamic mapping that reads date, datetime, or smalldatetime fields from a source and creates an Azure SQL Data Warehouse target at run time, the mapping fails to run on the Databricks Spark engine.
OCON-17119	When you use the default values for field delimiter or quote character in advance source or target properties, mappings fail with incompatible data types or precision errors or the Data Integration Service writes incorrect data to the target.

Bug	Description
OCON-12973	When you run a mapping that contains unconnected ports on the Hive engine, the mapping might fail or data corruption might happen.
OCON-12844	When you upsert or update data to Microsoft Azure SQL Data Warehouse and more than one column in the source table contains same value as the target column on which the primary key is defined, the Data Integration Service updates data incorrectly.

PowerExchange for Microsoft Azure SQL Data Warehouse Third-Party Limitations (10.4.0)

The following table describes third-party known limitations:

Bug	Description
OCON-21693	If the Informatica domain is SSL-enabled, dynamic mappings fail in the native environment. Microsoft ticket reference number: 119081423001742

PowerExchange for Microsoft Dynamics CRM

PowerExchange for Microsoft Dynamics CRM (10.4.1)

There are no fixed or known limitations in this release.

Review the Release Notes of previous releases for information about previous fixed limitations.

PowerExchange for Microsoft Dynamics CRM (10.4.0)

There are no fixed or known limitations in this release.

PowerExchange for MongoDB

PowerExchange for MongoDB (10.4.1)

There are no fixed or known limitations in this release.

Review the Release Notes of previous releases for information about previous fixed limitations.

PowerExchange for MongoDB (10.4.0)

PowerExchange for MongoDB Fixed Limitations (10.4.0)

There are no fixed limitations in this release.

Review the Release Notes of previous releases for information about previous fixed limitations.

PowerExchange for MongoDB Known Limitations (10.4.0)

The following table describes known limitations:

Bug	Description
OCON-11546	You cannot run Mongo shell commands from the Developer Tool. (405297)
OCON-11383	The performance of inserting data to MongoDB with 3.0 WiredTiger storage engine is slow when compared to MMapV1 storage engine. (433440)
OCON-11289	Transformation errors appear in the session log even if the mapping does not contain any transformations. Workaround: You can ignore the transformation errors. (406337)
OCON-11234	You cannot use the MongoDB ODBC driver to import a MongoDB collection that does not contain data. (405299)
OCON-11159	Column names of a MongoDB data object cannot exceed 255 characters. (344934)
OCON-11151	When you refresh a MongoDB data source in the Import Tables dialog box, the dropped MongoDB collections still appear. (404363)
OCON-11135	The virtual tables appear in the Import Tables dialog box only when you import the table metadata. (404362)
OCON-9726	When you create a mapping and specify special characters or white spaces in the source or target name, the mapping fails.

PowerExchange for MongoDB Third-Party Limitations (10.4.0)

The following table describes third-party known limitations:

Bug	Description
PLAT-14869	When you use PowerExchange for MongoDB to insert, update, or delete records in the MongoDB database, the performance for each of these operations is slow for a MongoDB cluster when compared to a single node. (437855)
OCON-11418	When you enable reading or writing as JSON documents, JSON import might fail because the JSON import and export formats in the Informatica MongoDB driver are not symmetrical. Workaround: To load JSON documents, use the mongoimport tool. (408233)
OCON-11231	Even when you enable Show array count in virtual main table , the Informatica MongoDB driver does not create the array count column. (405302)
OCON-11175	When you run a mapping with a Joiner transformation to read from a MongoDB database prior to version 2.6.3, the following error might appear in the mapping log even if the nested element is not null: LEFT_SUBFIELD only supports Object. (407700)

PowerExchange for Netezza

PowerExchange for Netezza (10.4.1)

There are no fixed or known limitations in this release.

Review the Release Notes of previous releases for information about previous fixed limitations.

PowerExchange for Netezza (10.4.0)

PowerExchange for Netezza Fixed Limitations (10.4.0)

There are no fixed limitations in this release.

Review the Release Notes of previous releases for information about previous fixed limitations.

PowerExchange for Netezza Known Limitations (10.4.0)

The following table describes known limitations:

Bug	Description
OCON-9071	When you run a mapping on the Blaze engine to read data that contains time or timestamp data types from a Netezza source, the data gets corrupted.
OCON-1111	When you create a Netezza connection by using the infacmd isp CreateConnection command, you must enter even the optional fields.
OCON-936	In Informatica Administrator, while deploying an application to a Data Integration Service, the physical data object type is incorrectly displayed for the Netezza data object.
OCON-658	When you read data from two or more Netezza sources, you cannot override the source schema and source table name at run time.
OCON-401	When run specify an incorrect table name in the Native Name field of a Netezza data object read operation, the mapping does not fail.

PowerExchange for OData

PowerExchange for OData (10.4.1)

There are no fixed or known limitations in this release.

Review the Release Notes of previous releases for information about previous fixed limitations.

PowerExchange for OData (10.4.0)

PowerExchange for OData Fixed Limitations (10.4.0)

There are no fixed limitations in this release.

PowerExchange for OData Known Limitations (10.4.0)

The following table describes known limitations:

Bug	Description
OCON-11505	OData service requests do not use the proxy server that is configured for the Data Integration Service. (417463)
OCON-1196	When the data size is greater than 700 MB, the OData mappings fail. (424279)

PowerExchange for Salesforce

PowerExchange for Salesforce (10.4.1)

PowerExchange for Salesforce Fixed Limitations (10.4.1)

Review the Release Notes of previous releases for information about previous fixed limitations.

The following table describes fixed limitations:

Bug	Description
OCON-16605	PowerExchange for Salesforce does not support dynamic mapping.

PowerExchange for Salesforce Known Limitations (10.4.1)

There are no known limitations in this release.

PowerExchange for Salesforce (10.4.0)

PowerExchange for Salesforce Fixed Limitations (10.4.0)

There are no fixed limitations in this release.

Review the Release Notes of previous releases for information about previous fixed limitations.

PowerExchange for Salesforce Known Limitations (10.4.0)

The following table describes known limitations:

Bug	Description
OCON-17050	PowerExchange for Salesforce does not support PowerCenter mapplets with Salesforce Merge, PickList, and Lookup transformation.
OCON-16605	PowerExchange for Salesforce does not support dynamic mapping.

Bug	Description
OCON-9108	The Data Integration Service does not create success files and error files for Salesforce mappings.
OCON-8948	When you create a mapping to write data from a Salesforce object to a database object and set CDC Time Limit to -1 to capture changed data for an infinite period of time, the mapping fails to load data into the target. This error occurs when the data read from the source is less than the batch size.

PowerExchange for Salesforce Marketing Cloud

PowerExchange for Salesforce Marketing Cloud (10.4.1)

There are no fixed or known limitations in this release.

Review the Release Notes of previous releases for information about previous fixed limitations.

PowerExchange for Salesforce Marketing Cloud (10.4.0)

There are no known limitations in this release.

PowerExchange for SAP NetWeaver

PowerExchange for SAP NetWeaver (10.4.1)

PowerExchange for SAP NetWeaver Fixed Limitations (10.4.1)

Review the Release Notes of previous releases for information about previous fixed limitations.

The following table describes fixed limitations:

Bug	Description
OCON-25420	When you write a mapping to write data with the timestamp data type to SAP BW version 7X, the Data Integration Service attempts to map the timestamp with the SAP DATS data type resulting in a data mismatch and the SAP BW mapping fails.
OCON-24925	When you read data that contains null values in the date fields from SAP BW using an Open Hub Destination (OHD), the Data Integration Service incorrectly interprets the null values as 1970-01-01.

PowerExchange for SAP NetWeaver Known Limitations (10.4.1)

There are no known limitations in this release.

PowerExchange for SAP NetWeaver (10.4.0)

PowerExchange for SAP NetWeaver Fixed Limitations (10.4.0)

There are no fixed limitations in this release.

PowerExchange for SAP NetWeaver Known Limitations (10.4.0)

The following table describes known limitations:

Bug	Description
OCON-19755	You cannot join SAP pool tables with cluster tables.
OCON-11528	In an SAP Table data object read operation, when you define a sort condition for cluster and pool tables and run a mapping, the mapping fails. (428266)
OCON-11497	When you add an SAP Table data object read operation as a lookup in a mapping, and configure the lookup condition based on a transparent table and a cluster or pool table, the mapping fails. (432043)
OCON-11454	When you enter FTP and SFTP details in an SAP connection and then clear the Use FTP and Use SFTP options, the details that you entered are cleared from the dialog box. This issue occurs when you create an SAP connection in Informatica Administrator and Informatica Developer. (407606)
OCON-11286	<p>On Windows 64-bit operating systems, when you run a mapping with the deprecated SAP data object read operation, the mapping fails.</p> <p>Workaround: Download the following library files from the SAP Service Marketplace:</p> <ul style="list-style-type: none">- icudt34.dll- icuin34.dll- icuuc34.dll <p>Copy the files to the following directory and run the mapping again:</p> <pre><Informatica installation directory>/services/shared/bin</pre> <p>To test SAP connections from Informatica Administrator, you must also copy the files to the following directory:</p> <pre><Informatica installation directory>/server/bin</pre> <p>(441761)</p>
OCON-11222	In Informatica Administrator, after you create a successful SNC-enabled SAP connection, if you edit the connection and specify incorrect values for the SNC library path, the Administrator tool validates the connection as successful. It does not display any error. (407551)
OCON-9961	If you create an SAP BW OHS mapping or workflow within a folder and do not specify the folder name when you set the OHD parameters, the Data Integration Service does not read any data. It also prints an incorrect OHS request ID in the mapping log.
OCON-9717	The Developer tool does not validate the values that you enter in the FTP fields of an SAP connection.
OCON-520	When you parameterize an SAP Table data object read operation property, you cannot define a precision that is higher than 28 for decimal data types. (435156)

PowerExchange for SAP NetWeaver Third-Party Limitations (10.4.0)

The following table describes third-party known limitations:

Bug	Description
OCON-10144	If you configure tracing in the <code>sapnwrfc.ini</code> file and use patch 42 of the SAP NetWeaver RFC SDK libraries, PowerExchange for SAP NetWeaver does not generate the trace file. SAP ticket reference number: 370381/2017 Workaround: Use patch 20 or patch 38 of the SAP NetWeaver RFC SDK libraries.

PowerExchange for Snowflake

PowerExchange for Snowflake (10.4.1)

PowerExchange for Snowflake Fixed Limitations (10.4.1)

There are no fixed limitations in this release.

PowerExchange for Snowflake Known Limitations (10.4.1)

The following table describes known limitations:

CR	Description
OCON-25386	The source pushdown optimization is not supported for the sorter transformation in a mapping.
OCON-24679	You cannot use external tables as Snowflake sources in mappings that run on the Spark engine.

PowerExchange for Snowflake Third-Party Known Limitations (10.4.1)

The following table describes third-party known limitations:

Bug	Description
OCON-24920	<p>When you run a Snowflake mapping on an AIX machine, the mapping fails with the following error:</p> <pre>net.snowflake.client.core.SFStatement executeQueryInternal SEVERE: Exception creating result java.lang.ExceptionInInitializerError at java.lang.J9VMInternals.ensureError(J9VMInternals.java:134)</pre> <p>Workaround: Set the following parameter in the Additional JDBC URL Parameters field in the Snowflake connection properties:</p> <pre>jdbc_query_result_format=json</pre> <p>Snowflake reference number: 00113748</p>

PowerExchange for Snowflake (10.4.0)

PowerExchange for Snowflake Fixed Limitations (10.4.0)

Review the Release Notes of previous releases for information about previous fixed limitations.

The following table describes fixed limitations:

Bug	Description
OCON-17057	When the Snowflake source contains case-sensitive table names and column names, the mapping fails on the Spark engine.

PowerExchange for Snowflake Third-Party Fixed Limitations (10.4.0)

Review the Release Notes of previous releases for information about previous fixed limitations.

The following table describes fixed limitations:

Bug	Description
OCON-16276	When two Snowflake source tables contain the same column name, and you specify a Joiner condition in the Snowflake source properties, the mapping fails on the Spark engine.

PowerExchange for Snowflake Known Limitations (10.4.0)

The following table describes known limitations:

CR	Description
OCON-23260	When you configure multiple pre-SQL or post-SQL queries in a mapping and if there is a space after the last query that ends with the ; character, the mapping fails.
OCON-21723	When you configure pushdown optimization in a mapping to write data to Snowflake, the update as update, update else insert, and delete operations are not supported.

CR	Description
OCON-22177	When you run a mapping on the Spark engine to write data to a Snowflake target and the target table and column names are not in uppercase letters, the mapping fails.
OCON-16608	When you run a mapping on the Spark engine to write data to a Snowflake target, and the data contains the Binary data type, the mapping fails.

PowerExchange for Snowflake Third-Party Known Limitations (10.4.0)

The following table describes third-party known limitations:

Bug	Description
OCON-21946	If the Informatica domain is enabled with SSL and you configure a mapping to read from or write to Snowflake, the dynamic ports are not generated. Snowflake reference number: 00034954
OCON-17317	When you use Snowflake jars, version 2.11-2.4.3, the performance of the read operation of a Snowflake mapping on the Spark engine drops by 1.5X when compared to Snowflake jars of version 2.11-2.3.1. Snowflake reference number: 00034954
OCON-12175	When you run a mapping in the native environment to read large volumes of data from Snowflake, the mapping fails with an out of memory error. Workaround: Increase the Java heap memory size in the Data Integration Service properties, and restart the Data Integration Service. Snowflake reference number: 00028631
OCON-11651	When you run a mapping on the Spark engine to write data of the Time data type, the Data Integration Service does not write the data to the Snowflake table even though the mapping runs successfully. Snowflake reference number:00033380

PowerExchange for Tableau

PowerExchange for Tableau (10.4.1)

There are no fixed or known limitations in this release.

Review the Release Notes of previous releases for information about previous fixed limitations.

PowerExchange for Tableau (10.4.0)

PowerExchange for Tableau Fixed Limitations (10.4.0)

Review the Release Notes of previous releases for information about previous fixed limitations.

The following table describes fixed limitations:

Bug	Description
OCON-20551	When the number of projects in Tableau exceed 100, the mapping fails.

PowerExchange for Tableau Known Limitations (10.4.0)

There are no known limitations in this release.

PowerExchange for Tableau V3

PowerExchange for Tableau V3 (10.4.1)

There are no fixed or known limitations in this release.

Review the Release Notes of previous releases for information about previous fixed limitations.

PowerExchange for Tableau V3 (10.4.0)

PowerExchange for Tableau V3 Fixed Limitations (10.4.0)

Review the Release Notes of previous releases for information about previous fixed limitations.

The following table describes fixed limitations:

Bug	Description
OCON-20551	When the number of projects in Tableau exceed 100, the mapping fails.

PowerExchange for Tableau V3 Known Limitations (10.4.0)

There are no known limitations in this release.

PowerExchange for Teradata Parallel Transporter API

PowerExchange for Teradata Parallel Transporter API (10.4.1)

There are no fixed or known limitations in this release.

Review the Release Notes of previous releases for information about previous fixed limitations.

PowerExchange for Teradata Parallel Transporter API (10.4.0)

PowerExchange for Teradata Parallel Transporter API Fixed Limitations (10.4.0)

There are no fixed limitations in this release.

Review the Release Notes of previous releases for information about previous fixed limitations.

PowerExchange for Teradata Parallel Transporter API Known Limitations (10.4.0)

The following table describes known limitations:

Bug	Description
OCON-1197	When you apply a filter condition on multiple ports either in the filter transformation or in the filter property of a Teradata source object, mapping fails on the Blaze engine.
OCON-658	When you read data from two or more Teradata sources, you cannot override the source schema and source table name at run time.
OCON-554	When you create a Teradata Parallel Transporter API connection by using the infacmd isp CreateConnection command, you must enter even the optional fields.

PowerExchange for Twitter

PowerExchange for Twitter (10.4.0)

There are no fixed or known limitations in this release.

Review the Release Notes of previous releases for information about previous fixed limitations.

PowerExchange for Web Content-Kapow Katalyst

PowerExchange for Web Content-Kapow Katalyst (10.4.0)

There are no fixed or known limitations in this release.

Review the Release Notes of previous releases for information about previous fixed limitations.

Informatica Global Customer Support

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

To find your local Informatica Global Customer Support telephone number, visit the Informatica website at the following link:

<https://www.informatica.com/services-and-training/customer-success-services/contact-us.html>.

To find online support resources on the Informatica Network, visit <https://network.informatica.com> and select the eSupport option.