



Informatica 10.2.2 Service Pack 1 Big Data Release Notes May 2019

© Copyright Informatica LLC 1998, 2020

Contents

Informatica 10.2.2 Service Pack 1.....	2
Support Changes.....	3
Deferred Support.....	3
Technical Preview Support.....	3
Verify System Requirements.....	3
Upgrade Path.....	4
Installation.....	4
Service Pack Files.....	4
Download the Service Pack Files.....	4
Update the Input.properties File.....	5
Run the Installer.....	6
Post-Installation Steps.....	6
Post-Installation Steps for the Analyst Service.....	6
Post-Installation Step for Enterprise Data Catalog.....	6
Post-Installation Steps for the Reference Data Database.....	6
10.2.2 Service Pack 1 Fixed Limitations.....	8
Big Data Management Fixed Limitations.....	8
Big Data Streaming Fixed Limitations.....	8
Big Data Quality Fixed Limitations.....	9
Enterprise Data Catalog Fixed Limitations.....	9
10.2.2 Service Pack 1 Known Limitations.....	9
Application Service Known Limitations.....	9
Big Data Management Known Limitations.....	10
Big Data Streaming Known Limitations.....	11
Enterprise Data Catalog Known Limitations.....	12
Enterprise Data Preparation Known Limitations.....	12
Third-Party Known Limitations.....	12
Transformations Known Limitations.....	14
10.2.2 Fixed Limitations and Closed Enhancements.....	14
Application Service Fixed Limitations.....	14

Big Data Management Fixed Limitations.	14
Developer Tool Fixed Limitations.	17
Enterprise Data Catalog Fixed Limitations and Closed Enhancements.	17
Enterprise Data Lake Fixed Limitations.	22
Mappings and Workflows Fixed Limitations and Closed Enhancements.	23
Parameters Fixed Limitations.	24
Profiles and Scorecards Fixed Limitations.	24
Security Fixed Limitations.	25
Third-Party Fixed Limitations.	25
10.2.2 Known Limitations.	26
Application Service Known Limitations.	26
Big Data Management Known Limitations.	26
Big Data Streaming Known Limitations.	30
Developer Tool Known Limitations.	30
Enterprise Data Catalog Known Limitations.	31
Enterprise Data Lake Known Limitations.	35
Profiles and Scorecards Known Limitations.	35
Reference Data Known Limitations.	36
Rule Specifications Known Limitations.	37
Third-Party Known Limitations.	37
Transformations Known Limitations.	40
Workflows Known Limitations.	41
Emergency Bug Fixes Merged into 10.2.2 Service Pack 1.	41
Informatica Global Customer Support.	41

This document contains important information about restricted functionality, known limitations, and bug fixes in Informatica 10.2.2 Service Pack 1.

Informatica 10.2.2 Service Pack 1

Informatica 10.2.2 Service Pack 1 contains multiple emergency bug fixes. The service pack supports Big Data Management, Big Data Streaming, Enterprise Data Catalog, Enterprise Data Preparation, and Big Data Quality.

To get the service pack, contact Informatica Global Customer Support.

Support Changes

This section describes the support changes in version 10.2.2 Service Pack 1.

Deferred Support

Effective in version 10.2.2 Service Pack 1, Informatica deferred support for the following functionality:

Streaming for the Data Masking Transformation

With the changes in streaming support, Data Masking transformation in streaming mappings is deferred.

Informatica intends to reinstate it in an upcoming release, but might choose not to in accordance with changing market or technical circumstances.

Technical Preview Support

Effective in version 10.2.2 Service Pack 1, Informatica includes the following functionality for technical preview:

SSL enabled Kafka connections for streaming

When you create a Kafka connection, you can use SSL connection properties to configure the Kafka broker.

Support for Hortonworks HDP 3.1

You can use Informatica big data products with Hortonworks HDP version 3.1.

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.

Verify System Requirements

Verify that your environment meets the minimum system requirements, such as operating systems and Hadoop distributions.

To see a list of the latest supported versions, see the Product Availability Matrix on the Informatica network:

<https://network.informatica.com/community/informatica-network/product-availability-matrices>.

Note: Upgrade all Informatica services and repositories running in a domain to Informatica 10.2.2 before you apply Informatica 10.2.2 Service Pack 1.

Upgrade Path

Install 10.2.2 Service Pack 1 on Informatica 10.2.2. Download and install the service pack.

Note: For information about Support EOL statements, contact Informatica Global Customer Support or see <https://network.informatica.com/docs/DOC-16182>.

Installation

To install the service pack, you can download the installer files, update the *Input.properties* file, and run the installer.

You can run the installer to apply or to roll back all product components or just a product component. After you install or roll back the service pack, complete the post-installation steps.

To roll back the service pack, set the rollback property in the *Input.properties* file, and then run the installer.

Service Pack Files

Informatica provides the service pack in .tar file and .zip file formats. After you download the service pack, extract the file contents. The service pack is available for Linux and Windows installations.

The service pack includes the following files:

Input.properties

Identifies the root directory of the Informatica installation to which you will install the service pack.

You update the file with the directory path. The file also contains a rollback property that you can set if you decide to uninstall the service pack.

install.bat

Installs the service pack to the directory that you specify on a Windows machine. Find the file in the Windows installer.

install.sh

Installs the service pack to the directory that you specify on a Linux machine. Find the file in the Linux installer.

Download the Service Pack Files

Download one or more installer files to install a service pack or roll back a service pack after you install.

To apply the service pack for 10.2.2 Service Pack 1, you can download the server installer, Informatica Developer installer, and command line utilities package.

Download one of the following packages:

- **informatica_1022SP1_server_linux-x64.tar**

Contains updates for Redhat Enterprise Linux installations.

- **informatica_1022SP1_server_suse11-x64.tar**

Contains updates for SUSE Linux Enterprise Server installations.

Download the following package for Informatica Developer:

- **informatica_1022SP1_client_winem-64t.zip**

Contains updates for Windows Developer tool installation.

Download one of the following packages for command line utilities:

- **informatica_1022SP1_cmd_utilities_linux-x64.zip**

Contains updates for Redhat Enterprise Linux installations.

- **informatica_1022SP1_cmd_utilities_suse11-x64.zip**

Contains updates for SUSE Linux Enterprise Server installations.

Update the Input.properties File

The *Input.properties* file includes properties that identify the Informatica installation and define the action taken when you run the service pack installer file. Update the properties before you install or roll back the service pack. Update the file in each service pack that you download.

1. Extract the service pack file.
2. Find the *Input.properties* file in the service pack.
3. Update the *DEST_DIR* property in the file with the path to the Informatica root directory.
 - On a Linux machine, set the path in the following format:
`DEST_DIR=/home/infauser/<version number>`
 - On a Windows machine, set the path in the following format:
`DEST_DIR=C:\\Informatica\\<version number>`
4. Configure the value of the *ROLLBACK* property. You can apply or roll back the service pack for all product components or to specific component. To install the service pack, retain the default value of **0**. To roll back the service pack, set the value to **1**.

When you install or roll back the service pack, the installer applies all the components, by default. To apply or roll back a specific component, remove the comment tag (#) associated with the component that you want to apply. For Big Data Management component, set *BDM_ONLY* to 1. For Enterprise Data Catalog, set *EDC_ONLY* to 1. For Enterprise Data Preparation, set *EDP_ONLY* to 1.

You must apply individual components for a customized product application in the following order:

1. Big Data Management
2. Enterprise Data Catalog
3. Enterprise Data Preparation

You must roll back the components for a customized product application in the following order:

1. Enterprise Data Preparation
 2. Enterprise Data Catalog
 3. Big Data Management
5. Save and close the file.

Run the Installer

Run the installer file to install the service pack or roll back the service pack after you install.

1. Close all Informatica applications and stop all Informatica services.
2. Find the installer file in the service pack files and extract the file.
 - For Linux systems, the installer file is *install.sh*.
 - For Windows systems, the installer file is *install.bat*.

Post-Installation Steps

After you install the service pack, perform the post-installation tasks that apply to your product.

Post-Installation Steps for the Analyst Service

After you download and apply the service pack, perform the following steps:

1. Verify that the Analyst Service is not running.
2. Delete the following directories from the 10.2.2 installation:
 - <Informatica root directory>/services/AnalystService/analyst
 - <Informatica root directory>/services/AnalystService/analystTool
 - <Informatica root directory>/services/AnalystService/mappingspec
 - <Informatica root directory>/tomcat/temp/<analyst_service_name>
If the *temp* directory contains multiple Analyst Service directories, delete the directory for each Analyst Service.
3. Restart the Analyst Service.
4. Clear the browser cache on the client machines.

Post-Installation Step for Enterprise Data Catalog

To implement the fix included in the emergency bug fix EBF-14316, bundled with 10.2.2 SP1, make sure that you index the Catalog Service after you upgrade to Enterprise Data Catalog 10.2.2 SP1.

Post-Installation Steps for the Reference Data Database

To run mappings that use reference tables on the Spark engine, perform the following steps after you install or upgrade:

1. Download the files that support JDBC connections for the reference data database.
2. Configure custom properties on the Data Integration Service.

Note: To find the reference data database that the Data Integration Service uses, examine the Content Management Service that the domain associates with the Data Integration Service. The Content Management Service identifies the reference data database connection. (BDM-17909)

Download the JDBC .JAR Files

Download the files that support JDBC connections on the reference data database. Copy the files to the Informatica services machine.

1. Obtain the JDBC .jar files that correspond to the reference data database that you use.
You can download the files from the database vendor web site.
2. Copy the file that you download to the following location:
<INFA_HOME>/externaljdbcjars

Configure Custom Properties on the Data Integration Service

Configure the custom properties in the Data Integration Service. Add the custom property that relates to the database type of the reference data database.

1. Open the Administrator tool and select the Data Integration Service in the Domain Navigator.
2. Find the Custom Properties option on the **Properties** tab.
3. Open the **Edit Custom Properties** dialog box.
4. Create a custom property for the reference data database.

The following table describes the custom properties that you can set for each database type:

Database Type	Property Name	Property Value
IBM DB2	ExecutionContextOptions.SparkRefTableHadoopConnectorArgs	--driver com.ibm.db2.jcc.DB2Driver --connect jdbc:db2://[db_hostname]:[port]/[database_name]
Microsoft SQL Server	ExecutionContextOptions.SparkRefTableHadoopConnectorArgs	--connect jdbc:sqlserver://[db_hostname];databaseName=[database_name]
Oracle	ExecutionContextOptions.SparkRefTableHadoopConnectorArgs	--connect jdbc:oracle:thin:@[db_hostname]:[port]:[sid]

5. Recycle the Data Integration Service.

10.2.2 Service Pack 1 Fixed Limitations

Big Data Management Fixed Limitations

The following table describes known limitations:

Bug	Description
PLAT-23580	A user in a group loses operating system profile privileges when you synchronize the user and the group through LDAP.
OCON-19308	When you run a mapping to read DATS data from an SAP Table source, the mapping fails with the following error: An error occurred while converting the value of the field [DATE_FROM] of type [DATS] to type [date/time]
OCON-18427	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: ERROR AbstractGoogleAsyncWriteChannel: Exception not convertible into handled response java.io.IOException: Write end dead
BDM-24853	Failed to recycle the REST Operations Hub Service without sending any REST request while the service is available.
BDM-24592	The Monitoring tool does not display the workflow graph for a workflow that you select.
BDM-24307	When the Blaze engine processes timestamp values from Parquet files on Hive, the engine truncates the timestamp values and does not write the nanoseconds to the target.

Big Data Streaming Fixed Limitations

The following table describes fixed limitations:

Bug	Description
IIS-2448	When you run a mapping with Joiner transformation, the output is not as expected because the late data is not cleared even after watermark delay.

Big Data Quality Fixed Limitations

The following table describes known limitations:

Bug	Description
IDQ-7917	When you test a large rule specification, the Analyst tool might become unresponsive. The issue can arise when the rule specification contains more than twelve rule sets.
IDQ-7863	The Analyst tool can take several minutes to compile a rule specification for the first time.
IDQ-7801	A mapping that contains a Standardizer transformation fails to run when the following conditions are true: <ul style="list-style-type: none">- The name of a strategy on the transformation includes one or more special characters.- The mapping originates in an Informatica 9.6.1 Model repository that you upgraded Informatica 10.2.

Enterprise Data Catalog Fixed Limitations

The following table describes known limitations:

Bug	Description
LDM-7416	The Informatica Platform resource fails with the following error when the JAVA_TOOL_OPTIONS environment variable is set in the cluster node: "Invalid command option: [xmx1g]".
EIC-16115	When you perform data discovery on JSON files, the profile run fails and exceptions appear if a column in the file has NULL data type.
EIC-16153	Enterprise Data Catalog takes more than 20 hours to export a Hive resource with one million assets.
EIC-16781	Deployment of Enterprise Data Catalog on a Cloudera cluster enabled for SSL and Kerberos with LDAP authentication fails with the following error: The difference in the Cluster's Core site and CCO core site regarding the property "hadoop.security.group.mapping.ldap.ssl.keystore.password". Cluster has "localjceks://file/CMF_CONF_DIR/creds.localjceks" (Parameterized), whereas CCO has complete path.

10.2.2 Service Pack 1 Known Limitations

Application Service Known Limitations

Bug	Description
OCON-19438	If the Metadata Access Service is inactive for a long time and you import a Hive, HBase, MapR-DB, or complex file data object, the import might fail with the following error: <pre>java.util.concurrent.RejectedExecutionException</pre> Workaround: Restart the Metadata Access Service and import the object again.

Big Data Management Known Limitations

The following table describes known limitations:

Bug	Description
BDM-24961	<p>A mapping that accesses Hive sources and targets and which stages temporary results in the staging directory fails with an error like:</p> <pre>SEVERE: [GRIDDTM_1016] The Integration Service failed to execute grid mapping with following error [An internal exception occurred with message: java.lang.RuntimeException: Failure to execute Query CREATE EXTERNAL TABLE</pre> <p>Examples of mappings that stage temporary results are mappings with a SQL override query to a Hive source, mappings that contain an Update Strategy Transformation, or mappings that create a transactional table.</p> <p>The mapping fails when you run it using the Blaze engine, or when you run it using the Spark engine on HDP 3.1. If you run it using the Spark engine on HDP 2.6, it succeeds.</p> <p>This issue occurs because the property <code>hive.server2.enable.doAs</code> is set to <code>FALSE</code> and thus the cluster is not configured to use user impersonation.</p> <p>Workaround:</p> <p>If the cluster uses Sentry or Ranger authorization, then verify that the Blaze working directory is managed by Sentry and Ranger, and verify that the impersonation user and the Hive user have full permissions on the Blaze working directory.</p> <p>If the cluster does not use Sentry or Ranger authorization, then edit <code>hive-site.xml</code> to set the property <code>hive.server2.enabled.doAs</code> to <code>TRUE</code>.</p>
BDM-24919	<p>The Blaze engine does not read source data with compression format from a Hive table when the following conditions are true:</p> <ul style="list-style-type: none"> - The table was created with a compression format. - Table data was not loaded through an Informatica mapping. <p>Workaround: Run mappings with an external hive table source on the Spark engine.</p>
BDM-24789	<p>If a mapping on the Blaze engine reads data from a Hive table and you specify database names in both the Data Access Connection String and in the runtime properties, a SQL override uses the database in the Hive connection instead of the database in the data object.</p>
BDM-24787	<p>If a mapping on the Spark engine reads data from a Hive table and you specify the database name in the data object, a SQL override uses the database in the Hive connection instead of the database in the data object.</p>
OCON-19697	<p>When you use a Hive connection to create a physical data object with a schema name which is different from the one specified in the Connection String, you cannot use a custom query to read data from a Hive source as the mapping results might be inconsistent.</p>
OCON-19687	<p>If you configure the <code>-DINFA_HADOOP_DIST_DIR</code> property to a value other than the default value in the <code>developerCore.ini</code> file, importing a complex file fails on the Developer tool machine.</p> <p>Workaround: Remove the value configured for the <code>-DINFA_HADOOP_DIST_DIR</code> property and reimport the object.</p>
OCON-19627	<p>A mapping can fail on the Spark engine with Amazon EMR version 5.2.0 when the following conditions are true:</p> <ul style="list-style-type: none"> - The mapping includes a transformation that reads one or more reference tables. - A column name in a reference table uses mixed-case characters.

Bug	Description
OCON-19518	A mapping can fail on the Spark engine with MapR version 6.1.0 when the following conditions are true: <ul style="list-style-type: none"> - The mapping includes a transformation that reads one or more reference tables. - A column name in a reference table uses mixed-case characters.
OCON-19467	When you select a Hive table from the Select a resource dialog box, the tables from other databases are also selected. Workaround: In the Connection Explorer, minimize other databases and expand only the database that you want to select.
OCON-19466	You cannot use the Show Default Schema Only option in the Connection Explorer to show tables that use the default schema and to add tables from non-default schemas.
OCON-19451	When you run a Hive mapping to read data from a custom data object and the source has two Hive tables with same name but from different databases, the mapping fails.
OCON-19312	When you run a Sqoop pass-through mapping on the Hortonworks HDP 3.1 distribution to write to a Hive target, the mapping is not optimized in the following scenarios: <ul style="list-style-type: none"> - You have enabled the Create or replace table at runtime target schema strategy and you do not specify a DDL query for create or replace table in the Hive target. - You have enabled the Retain Existing Target Schema option for a Hive target.
OCON-16421	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>

Big Data Streaming Known Limitations

The following table describes known limitations:

Bug	Description
IIS-3275	Rank transformation on a Kafka target topic with more than one partition produces unsorted data.
IIS-3242	The values assigned to the following JMS message header fields are ignored for JMS sources and targets because they are managed by JMS server: <ul style="list-style-type: none"> - JMSMessageID - JMSRedelivered - JMSTimestamp
IIS-3187	A streaming mapping stops responding when the <code>JMSCorrelationID</code> header field in a JMS message is set to byte array.
IIS-3185	The Data Integration Service writes data to any partition instead of to the Kafka partition specified through the Kafka partition ID header port.

Enterprise Data Catalog Known Limitations

The following table describes the known limitations:

Bug	Description
OCON-19561	When you run discovery on flat files that do not use UTF-8 encoding for character data, profiling fails and an exception appears.
EIC-17499	When you run discovery on XML files in Amazon S3 or File System resources, profile results do not appear for files that include a column name with special characters such as an underscore (_) or white spaces.
EIC-17168	If a column name in a resource ends with a space, the profiling scanner does not run the profile on the column.
EIC-17167	Unable to view a task on the Monitoring page in Catalog Administrator if there is no schedule configured for the task. The Catalog Administrator displays the <code>Schedule '[UNDEFINED]' not found</code> error message when you click the task.
EIC-17103	When you run the standalone scanner utility for a SQL Server Integration Services resource, test connection fails if the SSIS Scanner Type option is <i>File</i> .

Enterprise Data Preparation Known Limitations

The following table describes known limitations:

Bug	Description
IDL-11924	When you select a field of the MongoDB data type in a JSON file, the field appears to be empty.
IDL-12243	You cannot preview or publish an Avro file if the data lake is a Hadoop cluster that is secured with Apache Sentry.

Third-Party Known Limitations

The following table describes third-party known limitations:

Bug	Description
BDM-25070	The Spark engine fails to truncate external partitioned Hive tables on a Hortonworks HDP 3.1 cluster.
BDM-24903	When you run a mapping on the Spark engine with a partitioned Hive target, the mapping writes date/timestamp data in UTC format, resulting in incorrect data. The issue occurs when "Truncate target table" and "Truncate target partition" are enabled in the advanced properties of the mapping target object. For more information, see the following Hive issues: <ul style="list-style-type: none">- HIVE-21290- HIVE-21002

Bug	Description
BDM-24885	<p>When you run a Sqoop, Amazon S3, or Amazon Redshift mapping on the Blaze engine and on the Hortonworks HDP 3.1 cluster that uses Ranger KMS authorization, the mapping fails with the following error:</p> <pre>org.apache.hadoop.security.authentication.client.AuthenticationException: GSSException: No valid credentials provided (Mechanism level: Attempt to obtain new INITIATE credentials failed! (null))</pre> <p>Hortonworks HDP ticket reference number: 00220585</p>
BDM-24616	<p>A mapping that launches a HiveServer 2 task fails when it is configured to create a partitioned Hive target table with a char datatype column during runtime when the mapping runs on Hortonworks 3.1.</p> <p>Hortonworks HDP ticket reference number: 00219433</p>
BDM-24440	<p>A mapping that runs on a Cloudera CDH 6.x cluster cannot read data from a file that contains Unicode characters in the file name.</p> <p>Workaround: Set the LANG environment variable in the Hadoop connection.</p> <p>Cloudera CDH ticket reference number: 356638</p>
BDM-21486	<p>If a mapping runs on an Azure HDInsight cluster, the mapping fails when you use a SQL override on a Date/Time column in a Hive source.</p> <p>Apache Hive ticket reference number: HIVE-12200</p>
BDM-20346	<p>The Spark engine fails to run a built-in Hive SerDe of type CSV on a Hortonworks HDP 3.1 cluster.</p> <p>Hortonworks HDP ticket reference number: 00220586</p>
IDL-12225	<p>You cannot preview a JSON file if the data lake is on Hortonworks HDP 3.1.0.</p>
OCON-19673	<p>In a Cloudera CDH 6.1 environment, Parquet files cannot be imported into the Developer tool because arrays are represented as bag data types in Parquet files on Cloudera CDH 6.1.</p>
OCON-19610	<p>When you run a Hive mapping to write data to a Hive target in Avro file format on the Hortonworks HDP 3.1 distribution, the mapping fails with the following error:</p> <pre>Storage schema reading not supported</pre>
OCON-19506	<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, multiple Spark jobs fail due to transient errors in Google Cloud Storage. However, the mapping runs successfully with the following error:</p> <pre>java.io.IOException: java.io.IOException: Write end dead</pre>
OCON-19488	<p>When you use the Sqoop Hortonworks Connector for Teradata to run Teradata mappings on the Blaze or Spark engine with the Hortonworks HDP 3.1 cluster, the mapping fails with the following error:</p> <pre>NoClassDefFoundError: com/cloudera/sqoop/manager/ManagerFactoryNoClassDefFoundError</pre> <p>Workaround: To run a mapping on the Spark engine, download the 1.6.0 version of the TDCH driver from Cloudera: https://www.cloudera.com/downloads/connectors/sqoop/teradata/1-6c5.html</p> <p>Note: There is no workaround for the Blaze engine.</p>

Transformations Known Limitations

The following table describes known limitations:

Bug	Description
BDM-25112	<p>Mappings fail when the following conditions are true:</p> <ul style="list-style-type: none">- Mapping runs on the Spark engine.- The mapping contains a Union transformation.- The same physical data object source is used multiple times in the mapping. <p>You will see the following error:</p> <pre>java.lang.RuntimeException: java.io.IOException: No service instances found in registry</pre> <p>Workaround: Create a copy of the physical data object source to use in the mapping.</p>

10.2.2 Fixed Limitations and Closed Enhancements

Application Service Fixed Limitations

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

The following table describes fixed limitations:

Bug	Description
OCON-13454	The Developer tool fails to connect to an active Metadata Access Service process even if the Metadata Access Service is configured for high availability.
MRS-1795	When the Model Repository Service runs out memory, no Java heap dumps are generated.
MRS-1571	When you enable the Model Repository Service for DEBUG level of logging, the service generates multiple ERROR level log messages.
BDM-17402	When the Data Integration Service runs on a grid, the service waits for applications to load before the service starts up.
BDM-17301	The Data Integration Service restarts if deployed applications take more than 15 minutes to load during service start-up.
BDM-17293	Application details cannot be fetched from the Data Integration Service while the Data Integration Service is deploying applications.
BDM-17055	Requests to submit mappings hang if the application that contains the mapping is loading on the Data Integration Service during service start-up.

Big Data Management Fixed Limitations

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

The following table describes fixed limitations:

Bug	Description
OCON-17874	When you run a Sqoop mapping with a TLS-enabled domain and SparkSqoopDisSideInvocation property is set to true, the mapping fails because the location of the truststore system property of the Data Integration Service gets modified from <code>-Djavax.net.ssl.trustStore=\$INFA_HOME/services/shared/security/infa_truststore.jks</code> to <code>-Djavax.net.ssl.trustStore=\$INFA_HOME/externaljdbcjars/./java/jre/lib/security/cacerts</code> .
OCON-16535	When you run a Sqoop mapping on the Spark engine, the MapReduce job fails because the Data Integration Service ignores the <code>hadoop.security.token.service.use_ip</code> configuration property in the <code>core-site.xml</code> file.
IDL-8652	If you run a mass ingestion specification that contains more than 2,000 tables, a majority of the ingestion mapping jobs fail with the following error: <code>SEVERE: UserDefinedJob execution failed</code>
IDL-10354	When a run instance of a mass ingestion specification completes, the Summary page in the Mass Ingestion tool continues to display the status as RUNNING. The Summary page updates the status only after you navigate to the Execution Statistics page for the run instance.
BDM-23906	The <code>infacmd ccps listClusters</code> command fails when the domain contains a Databricks cluster that does not have any custom tags. Workaround: Manually add a custom tag to the cluster or create the cluster through a cluster workflow.
BDM-23599	Connectivity to Hive sources and targets on Amazon S3 with SSE-KMS authorization fails on a Hortonworks HDP 2.6.1 or higher version cluster with an error reading "The encryption method specified is not supported."
BDM-23199	When a sorted Joiner transformation joins data from a single source, data loss can occur if the row size for key columns exceed the data cache size.
BDM-22862	If the Spark engine runs a mapping that creates temporary Hive tables and the data contains string values with newline characters (<code>\n</code>), the newline character is treated as the end of one record. For example, if the temporary Hive table stores the string <code>'abc \n def'</code> , the first record is <code>'abc'</code> and the second record is <code>'def'</code> . One situation where the issue occurs is in a mapping that contains an Update Strategy transformation. When the mapping runs on the Spark engine, the Update Strategy transformation stages data in temporary Hive tables.
BDM-22754	When the DTM runs out of memory and shuts down, the subsequent mappings sometimes fail with a <code>NullPointerException</code> during mapping initialization.
BDM-22736	When you run a mapping with a large number of Joiner transformations on the Spark engine, the Data Integration Service runs out of memory.
BDM-21657	Blaze mapping reads NULL data from a Parquet format Hive table.
BDM-21324	Incorrect query generated for a Hive update strategy when the column name starts with a number.
BDM-18221	If the <code>tez.tez-ui.history-url.base</code> is set to NULL or if it is not present in the cluster configuration, the Tez jobs fail with a null pointer exception.

Bug	Description
BDM-18195	When you disable yarn.timeline-service in the cluster configuration, the Tez view does not show the Hadoop job.
BDM-18173	The duration taken to load the Summary Statistics and Execution Statistics in the Administrator tool for a non-Administrator user is more when compared to an Administrator.
BDM-17879	Slow performance of domain when OS profiles are enabled.
BDM-17849	When you use infacmd command line program to export a mapping that contains an XML payload of a complex data type, the export fails with the following error: <pre>com.sun.xml.xsom.parser.XSOMParser cannot be found by com.infa.products.sdk.adapter.parser.xml.utils_1.1.0 com/sun/xml/xsom/parser/XSOMParser [ICMD_10033] Command [ExportObjects] failed with error [[OIECMD_10006] Error occurred during export function. See previous error.]</pre>
BDM-17831	The JVM options for debugging are disabled for Data Integration Services.
BDM-17755	When a decimal port on a partitioned, bucketed Hive target is disconnected, a passthrough mapping that runs on the Blaze engine fails.
BDM-17698	When you use the Cloudera Connector Powered by Teradata to run a Sqoop TDCH mapping, the mapping fails with the following error: <code>Missing argument for option: staging-database</code>
BDM-17685	The Data Integration Service does not remove temporary empty files from the operating system profile target directory when all the following conditions are true: <ul style="list-style-type: none"> - The Data Integration Service is enabled to use operating system profiles. - The mapping writes to an HDFS target. - The mapping runs on the Hive engine.
BDM-17541	When you run a mapping by choosing Window > Preferences > Run configurations > Mapping , the Developer tool attempts to re-establish connections with all configured domains. Depending on the connection quality, this may result in dramatically slower performance.
BDM-17175	After you import the workflow into the Model repository, you cannot create a parameter set for the SQL parameter with the empty default value.
BDM-16843	After you import an SQL parameter into the Model repository, you can edit the parameter value but you cannot reset it later to the default empty or NULL value.
BDM-14626	When you do not assign a connection in the physical data object, the workflow application fails to deploy.
BDM-14139	In a Hortonworks HDP or an Amazon EMR environment, a mapping run on the Spark engine fails when all of the following conditions are true: <ul style="list-style-type: none"> - The mapping writes to an Avro target. - The schema contains an array or a map that exceeds two levels of nesting. - A struct element exists at the third or higher level of nesting. For example, when you write to an Avro target with a schema such as <code>array(array(struct))</code> or <code>array(map(string, struct))</code> , the mapping fails.
BDM-13107	When you run a mapping with Spark monitoring enabled, performance bottlenecks can appear. The performance varies based on the mapping complexity.

Developer Tool Fixed Limitations

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

The following table describes fixed limitations:

Bug	Description
OCON-14447	When you use an SQL transformation to read multiple records from a stored procedure, the data preview fails.
OCON-13057	When you use an ODBC connection with Microsoft SQL Server as the ODBC provider, and run a data preview on a Microsoft SQL Server named instance, data preview fails.
IDQ-7121	The Developer tool raises an error if you import a mapplet to the Model repository and the following conditions are true: <ul style="list-style-type: none">- The mapplet that you import replaces a mapplet in the repository.- The mapplet in the repository reads a token set that the replacement mapplet does not read.- The two mapplets have the same configuration, including a Labeler transformation that reads the token set.
BDM-20406	The Developer cannot open an object when there is a mismatch between the IDs of the local and the cached object. It returns the following error: <code>Duplicate object of ID</code>
BDM-18315	In the Developer tool, the Header command write property for flat files does not support any escape characters.
BDM-17541	When you run a mapping by choosing Window > Preferences > Run configurations > Mapping , the Developer tool attempts to re-establish connections with all configured domains. Depending on the connection quality, this may result in dramatically slower performance. To work around this issue, delete unused domain configurations from the Developer tool.

Enterprise Data Catalog Fixed Limitations and Closed Enhancements

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

The following table describes fixed limitations:

Bug	Description
LDM-7044	Sometimes, the test connection fails for an Axon resource if you do not select the Enable Secure Communication option when you create the resource.
LDM-5987	Enterprise Data Catalog and Catalog Administrator display junk characters in the search results and resource properties section after you change the server locale to Japanese.
EIC-8243	You cannot create a custom attribute with the same name as an existing system attribute.
EIC-8032	For an Apache Atlas instance running on a Kerberos-enabled Hortonworks Data Platform (HDP) cluster, the Apache Atlas resource extracts source and target information only for Hive.

Bug	Description
EIC-7969	When you perform an installation or an upgrade in the silent mode and specify a value other than 1 or leave the value blank for the property <code>INSTALL_LDM</code> , the Model Repository Service upgrade fails.
EIC-6422	<p>When you update the Catalog Service from any previous version to version 10.2.1 Update 1, the following error occurs:</p> <pre>ScannerPluginRegistrationManager@288] - [ModelValidationException_00016] Datatype for Attribute [FireonUpdated] has been altered. Expected: [core.Boolean]. Found: [core.Integer].</pre>
EIC-6384	After you import a CSV file that has assets with associated business titles, Enterprise Data Catalog does not display the associated business titles for the assets.
EIC-6080	In Enterprise Data Catalog, the Allow Filtering option for a default custom attribute named Department is not enabled even if you have enabled the Allow Filtering option for the custom attribute in the Catalog Administrator.
EIC-6073	Enterprise Data Catalog displays a mapping that is deleted even after you re-run the PowerCenter resource in Catalog Administrator.
EIC-5999	Enterprise Data Catalog displays incorrect search results when you search the catalog using an asset name that contains special characters.
EIC-5916	<p>An Amazon S3 resource fails with the following data validation error:</p> <pre>Error occured in scanner execution java.util.concurrent.ExecutionException: com.infa.products.ldm.filescanner.file.processor.api. FileProcessorException: java.util.concurrent.ExecutionException: com.infa.products.ldm.filescanner.file.processor.api. FileProcessorException: com.infa.products.ldm.ingestion.exceptions . DataValidationException: [DataValidationException_00009]</pre>
EIC-5906	When you initiate Catalog Service for Enterprise Data Catalog 10.2 Update 2, the Catalog Service does not start.

Bug	Description
EIC-5860	When you run a PowerCenter resource that has approximately 600 mappings with load type configured as medium and scanner memory set to high, the resource run fails with the error "String index out of range: -1".
EIC-5768	The lineage diagram and the schema name are not displayed in the popup that appears when you move the mouse pointer on an external object such as a table or a view.
EIC-5734	When you run a HDFS resource on an XML file that has more than 1000 elements, the profiling run fails for the resource.
EIC-5407	The Informatica Platform resource fails with the following error: "ArrayIndexOutOfBoundsException".
EIC-5374	Custom lineage for a data source that includes tables on either side of the seed asset displays the tables as a group.
EIC-5306	A resource fails when you renew the Kerberos ticket while the resource is running.
EIC-4640	Enterprise Data Catalog does not add a business title to a resource if the resource name contains a locale character.
EIC-4626	An Informatica Platform resource fails if a mapping contains unstructured files, such as PDF.
EIC-4545	You cannot delete multi-valued custom attribute values of reference data type in Enterprise Data Catalog.
EIC-4053	When you install Enterprise Data Catalog on an embedded cluster and you have a Web proxy server configured in <code>/etc/yum.conf</code> , you might encounter issues. To resolve this issue, the line <code>proxy=None</code> needs to be added to the following repository files: <ul style="list-style-type: none"> - ambari.repo - HDP.repo - HDP-utils.repo
EIC-2975	Duplicate mapping objects appear in the lineage and impact view for a table linked to a mapping containing the lookup transformation.
EIC-14407	The regular expression that selects the subset of a schema does not work for a Hive metadata resource.
EIC-13725	The import job fails with a null pointer exception when you update the custom attribute value of a reference data type in the import file. This issue occurs for business glossary assets.

Bug	Description
EIC-13723	When you run a resource, the Solr process waits for the .lock files in HDFS to be deleted.
EIC-13443	Sometimes, you cannot view the business title of a resource in Enterprise Data Catalog.
EIC-12993	The pre-validation utility validates 2 GB of free disk space required in the /usr and /var directories.
EIC-12822	For the Informatica Platform resource, additional mappings are extracted in Enterprise Data Catalog for the mapping instance when there are no mapping parameters.
EIC-12551	When you run profiling on a Salesforce resource, sometimes, the profiling fails for a few tables and the error "ERROR: nsort_define() returns -26" appears.
EIC-12493	The PowerCenter resource fails with the following error during the conversion of XML to XDoc files: "java.lang.ArrayIndexOutOfBoundsException".
EIC-12474	Enterprise Data Catalog displays the instance name instead of the database name for an Oracle resource.
EIC-12142	The pre-validation utility fails when it validates the prerequisites to create an Informatica Cluster Service. The failure occurs when you specify a custom Hadoop directory that is nested within another directory.
EIC-12057	Custom property added that you can use when you deploy Enterprise Data Catalog on a WANdisco Fusion-enabled cluster. You can use the property to resolve the issues on clusters where WANdisco Fusion is not enabled.
EIC-11880	The Ambari_server_launch.sh and the Ambari_agent_install.sh scripts add the export PYTHONHTTPSVERIFY=0 line in the /usr/sbin/ambari-server and the /usr/sbin/ambari-agent files every time you start the Informatica Cluster Service. The line must be added only if it does not exist.
EIC-11879	The Catalog Service uses a CPU core continuously even when there are no processes running or initiated.
EIC-11597	File System scanner fails when you use the Server Message Block (SMB) protocol.
EIC-11597	The File System resource fails when you use the Server Message Block (SMB) protocol to extract metadata from the file system.

Bug	Description
EIC-11484	After you import a CSV file that has assets with associated custom attribute values, Enterprise Data Catalog does not display the associated custom attribute values for the assets.
EIC-11371	<p>If the resource name, application name, and parameter file share the same name, the Informatica Platform resource fails with the following error:</p> <pre> ERROR com.infa.products.ldm.lineageservice. client.impl.LineageServiceClientImpl- Sep 14, 2018 3:21:46 PM com.infa.products.ldm.scanners.mrsscanner. cmd961bundle. LineageServiceCaller\$LineageServiceTask disconnect </pre>
EIC-11357	After you create and run an SAP Business Objects resource, Enterprise Data Catalog does not display the lineage information for the crystal reports.
EIC-11335	The lineage summary does not display flat files that are present in the Window File System resource. This issue occurs when you use FTP connection type to configure the Windows File System resource.
EIC-11289	When you create a Azure Microsoft SQL Server resource, enable data discovery, and run the resource, the temporary tables that are created during the resource run are not deleted by default.
EIC-11282	In the Overview tab, Enterprise Data Catalog does not display the source description of an associated business term for an asset. This issue occurs after you upgrade Enterprise Data Catalog from version 10.2.0 to 10.2.2.
EIC-11229	When you configure an Informatica Platform resource with an incorrect domain name, the resource fails and the scanner log file points to the lineage service log file incorrectly.
EIC-11220	If a custom attribute of reference data type includes the ampersand (&) symbol in the name, Enterprise Data Catalog appends "amp" with the custom attribute name.
EIC-11202	When you run an HDFS resource after you run a PowerCenter resource with Auto Connection Assignment, the HDFS resource run takes a long time to complete.
EIC-11178	When you run a resource that has empty source files, the profile run fails and a null pointer exception appears in the log file.

Bug	Description
EIC-11140	Data validation fails for an Erwin resource because the resource does not support multi-valued attributes.
EIC-11065	When an HDFS connection is parameterized, the connection assignment path displays the parameter value appending the path for a PowerCenter resource.
EIC-11053	<p>The test connection fails for the UCF.MongoDBNoSQLDatabase resource and the following error appears:</p> <pre>Error com.mongodb.MongoCommandException: Command failed with error 13: 'not authorized on admin to execute command { listDatabases: 1 }' on server inkrrh65qa01:27017. The full response is { "ok" : 0.0, "errmsg" : "not authorized on admin to execute command</pre>
EIC-11019	Sometimes, the inferred data domains appear in green instead of orange in the Similar Column section in the Overview tab.
EIC-11004	An incorrect URL is generated if you provide the region name in the Amazon S3 URL.
EIC-10255	Profiling fails for a resource if you add LDAP user instead of the domain user in the Data Integration Service properties of the resource.

The following table describes closed enhancement requests:

Bug	Description
EIC-5268	For the Informatica Platform resource, you cannot view the lineage in Enterprise Data Catalog when the custom SQL query is parameterized in the mapping.
EIC-13569	In Enterprise Data Catalog, you cannot view the lineage information for an asset after you assign a business title or custom attribute value to the asset.

Enterprise Data Lake Fixed Limitations

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

The following table describes fixed limitations:

Bug	Description
IDL-9988	When you perform operations such as import, publish, search, or data preview in Enterprise Data Lake, the application displays all Hive schemas in the data lake, including schemas for which you do not have permissions.
IDL-8888	When the name of the Informatica security domain to which your user account belongs contains spaces, and you perform an upload, publish, or download activity in Enterprise Data Lake, the activity is not displayed in the application home, activity, and project pages.

The following table describes closed enhancement requests:

Bug	Description
IDL-11420	The application no longer makes requests to external web sites during the log in process.

Mappings and Workflows Fixed Limitations and Closed Enhancements

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

The following table describes fixed limitations:

Bug	Description
MWF-1680	The Data Integration Service can leak memory when you run a large number of workflow applications and the workflows include mappings that run from an operating system profile.
MWF-1668	If you bind a workflow variable to a mapping output, the Data Integration Service does not update the variable value if the mapping output is null.
MWF-1639	If you bind the taskStatus output to an integer variable, a workflow task returns a taskStatus value of 0 regardless of the success or failure of the task.
MWF-1636	The Data Integration Service can generate an excessive quantity of threads and other resources when you run a workflow that sends an email, for example in a Notification task.
MWF-1631	If a workflow runs from an infacmd command and infacmd loses contact with the Data Integration Service, infacmd reports the workflow as failing to start while the workflow starts correctly.
MWF-1621	The workflow engine does not start if it encounters a locked process in the workflow database.
MWF-1597	You can save a workflow in which a task instance does not specify a task to perform, which can cause the workflow to fail at run time.
IDQ-7594	If you run multiple mappings with Match transformations concurrently, the mappings might generate multiple cache data files with the same name. The Data Integration Service might fail to delete all of the cache data files, and therefore the mappings might fail.
IDQ-7162	The Match transformation might not add all matching records to the correct clusters if the transformation input data includes very large record groups.

Bug	Description
BDM-20274	<p>If you run a mapping during the last week of the year and the week ends in the next year, the system parameter MappingRunStartTime increments the year by 1. The week begins on Sunday and ends on Saturday.</p> <p>For example, the last week of 2018 included Sunday, December 30th and Monday, December 31st. If you ran a mapping on December 30th or December 31st, the parameter had a value of 12/30/2019 or 12/31/2019 respectively.</p>
BDM-18771	The infacmd wfs upgradeWorkflowParameterFile command can take a long time to complete the upgrade of a set of workflow parameter files.
BDM-18115	<p>The LTRIM and RTRIM functions do not trim spaces when the following conditions are true:</p> <ul style="list-style-type: none"> - An expression is nested in multiple levels of LTRIM or RTRIM functions, such as <code>LTRIM(LTRIM(<expression>))</code> or <code>RTRIM(LTRIM(<expression>))</code>. - The nested expression references a column in a Hive table that originates from an SQL override.
BDM-17376	<p>You might get unexpected results in a mapping with a Union transformation when the following conditions are true:</p> <ul style="list-style-type: none"> - The Union transformation has at least three input groups. - The same two ports are unconnected in all the groups. - The mapping contains a downstream transformation that filters data. <p>The transformation might incorrectly filter data that should be retained, or it might retain data that should be filtered out.</p>

The following table describes closed enhancement requests:

Bug	Description
MWF-1638	The Data Integration Service writes log messages to describe workflow engine activity when the workflow engine starts up. The Data Integration Service writes the number of workflow instances that the workflow engine processes to the log. The Data Integration Service also writes the workflow name, application name, and instance ID to the log.

Parameters Fixed Limitations

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

The following table describes fixed limitations:

Bug	Description
BDM-17211	<p>If you resolve mapping parameters in a mapping that contains AppSDK data objects, an exception occurs when you perform the following actions:</p> <ul style="list-style-type: none"> - Run the mapping. - Run data preview. - Show the execution plan.

Profiles and Scorecards Fixed Limitations

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

The following table describes fixed limitations:

Bug	Description
IDE-3703	Enterprise discovery profile with primary key and foreign key inference options do not work when the profiling warehouse uses an ODBC connection with UseDSN option.
IDE-3640	In Informatica Analyst, when the Scorecard workspace is active and you perform any task in the Analyst tool, high CPU usage and memory utilization are observed.
IDE-3463	After you upgrade from 10.1.1 HotFix 1 to 10.2.1, you cannot run a few existing scorecards.
IDE-3448	You cannot run a scorecard after you delete a rule in the Metrics tab of the scorecard.
IDE-3396	After you upgrade to version 10.2, the scorecard run for logical data objects fails with null pointer exception.
IDE-3180	When you run a column profile on a source that has a column with the precision value greater than 255, you cannot infer data types on the column.
IDE-3161	In Informatica Analyst, when you use Internet Explorer to create a column profile that has a Japanese character in the profile name, the page becomes unresponsive.
IDE-3115	Column profile fails for a Hive partitioned source.
IDE-3065	When you run enterprise discovery profile with data domain discovery, profiling fails for some Hive sources.
BDM-19917	Profiling run fails when you run the <i>infacmd ps execute infacmd</i> command.
BDM-17440	When you create a column profile or data domain discovery profile on a partitioned Hive table and run the profile on the Blaze engine, the profile run intermittently fails.

Security Fixed Limitations

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

The following table describes fixed limitations:

Bug	Description
ANT-520	If a secure domain is SAML-enabled and you log in to the Administrator tool as an SAML user, the Administrator tool does not show the user activity log. The Analyst tool shows the user activity log but does not show the name of the user in the user field of the activity log.

Third-Party Fixed Limitations

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

The following table describes fixed limitations:

Bug	Description
BDM-10410	For Hortonworks clusters, insert into a Hive target table on Amazon S3 can fail if the table is enabled for transactions.

10.2.2 Known Limitations

Application Service Known Limitations

Bug	Description
OCON-13257	When you import a complex file object, an HBase object, or a Hive object from a Hadoop cluster, the metadata import fails if the domain and the Metadata Access Service use Kerberos authentication.
MRS-1617	When you configure the monitoring Model repository on an Oracle database, an exception appears when the number of objects fetched by the monitoring Model Repository Service for the mapping job statistics exceeds 1000 objects.
MRS-1585	An error appears when you enable Git version control system for a Model repository if the global repository name shares the same name with another Model repository, and if both the repositories share the local repository. Workaround: Use separate local repository for the Model Repository Services.
MRS-1570	After you assign the Show Security Details privilege, you cannot view upstream dependencies for a reusable Model repository object unless you have read permission on the dependent projects.
MRS-1463	The version history for Model Repository Service objects are lost if the content of the versioned Model repository is migrated from one Perforce depot to another.
MRS-1861	Sometimes, the out of memory error might appear for Model Repository Service. Workaround: Increase the Maximum Heap Size value in Advanced Properties to 1536m or 2048m.

Big Data Management Known Limitations

The following table describes known limitations:

Bug	Description
OCON-9377	When you configure Sqoop and run a Teradata Parallel Transporter mapping on a Cloudera cluster to export data of the Byte or Varbyte data type to a Teradata target, the mapping fails on the Blaze engine.
OCON-9376	If you configure Sqoop to export data of the Blob or Clob data type to a Teradata target, TDCH mappings fail on the Spark engine.

Bug	Description
OCON-9143	<p>In the read and write operations for a complex file data object, you cannot edit the precision and scale of elements within a field that is of a complex data type.</p> <p>For example, if Field1 is of type array with string elements, you cannot edit the precision and scale of the string elements.</p>
OCON-9005	<p>When you run TDCH mappings on the Hive engine to write time data to a Teradata target, the nanosecond part is truncated.</p>
OCON-8850	<p>If you configure Sqoop to export data of the Timestamp data type from a Hive source to a Microsoft Azure SQL Data Warehouse target, the mapping fails.</p>
OCON-8779	<p>If you configure Sqoop to export data of the Real data type to IBM DB2 z/OS targets, the mapping fails.</p>
OCON-7687	<p>When you export data through Sqoop and the columns contain mixed case characters, the mapping fails.</p>
OCON-7669	<p>When you configure Sqoop and OraOop, and export data to an Oracle target that contains mixed case characters in the table name, the mapping fails.</p> <p>Workaround: Use the generic Oracle JDBC driver to export data.</p>
OCON-7429	<p>When you run a Teradata Parallel Transporter mapping on a Hortonworks cluster and on the Blaze engine to write data of the Byte or Varbyte data type to a Teradata target, the data gets corrupted. This issue occurs when you use the <code>hdp-connector-for-teradata-1.5.1.2.5.0.0-1245-distro.tar.gz</code> JAR.</p> <p>Workaround: Use the <code>hdp-connector-for-teradata-1.4.1.2.3.2.0-2950-distro.tar.gz</code> JAR.</p>
OCON-730	<p>When you export data through Sqoop and there are primary key violations, the mapping fails and bad records are not written to the bad file. (456616)</p>
OCON-7216	<p>If a Sqoop source or target contains a column name with double quotes, the mapping fails on the Blaze engine. However, the Blaze Job Monitor incorrectly indicates that the mapping ran successfully and that rows were written into the target.</p>
OCON-7212	<p>If there are unconnected ports in a target, Sqoop mappings fail on the Blaze engine. This issue occurs when you run the Sqoop mapping on any cluster other than a Cloudera cluster.</p> <p>Workaround: Before you run the mapping, create a table in the target database with columns corresponding to the connected ports.</p>
OCON-7205	<p>When you run a Sqoop mapping on the Blaze engine to export data of the Numeric data type from Netezza, the scale part of the data is truncated.</p>
OCON-17245	<p>When you run a Sqoop mapping on a Kerberos-enabled HDInsight cluster with ADLS as storage, the mapping fails on the Blaze engine.</p>

Bug	Description
OCON-17194	<p>Microsoft Azure Blob Storage incorrectly appears as a source option for the Lookup transformation on the Spark and Databricks Spark engines, but this option is not supported. Mapping validation fails with the following error when you select Microsoft Azure Blob Storage as a lookup source on the Spark or Databricks Spark engine:</p> <pre>[Spark Validation] The Integration Service cannot use the Spark engine to push the transformation [test_input_txt_Read1] of type [LookupTx] to the Hadoop cluster. [ID: TRANSFORMATION_NOT_SUPPORTED] test_input_txt_Read1</pre>
OCON-16315	<p>The Data Integration Service fails with out of memory errors when you run a large number of concurrent mappings with data object read and write operations that project columns as complex data types.</p> <p>Workaround: Perform any of the following steps:</p> <ul style="list-style-type: none"> - Increase the heap memory settings on the machine where the Data Integration Service runs. - Reduce the number of concurrent mappings that process complex data types. - Set the Data Integration Service custom property ExecutionOptions.CBOStatCacheSize to a reasonably small number.
IN-3362	<p>A data object with an intelligent structure model might accept JSON input files formatted with newlines between attributes as suitable for partitioning, even though the files cannot be partitioned. This might cause data that corresponds to the model to be identified as lost data in the Spark log.</p>
BDM-4597	<p>A mapping with a joiner transformation that processes more than 4,294,967,294 rows in a single partition will fail.</p> <p>Workaround: If possible, increase partitioning on the source.</p>
BDM-23575	<p>The number of nodes is incorrectly displayed in the session log for the mappings that run on nodes that are labeled in a cluster that runs on Blaze and Spark engines.</p>
BDM-23550	<p>When an Update Strategy transformation contains an insert, update, or a delete operation, and a JDBC target, incorrect number of inserted, updated, or deleted rows appear from the Spark events.</p>
BDM-23392	<p>When a mapping that runs on Spark uses an Update Strategy transformation, table constraints might cause a BatchUpdateException and mapping failures.</p> <p>Workaround: Edit the Spark.JdbcNumPartition setting in the mapping Runtime Properties to reduce the number of partitions to 1. This causes all rows to be processed as a single partition.</p>
BDM-23317	<p>When you monitor statistics for the Data Integration Service configured with file-based queuing, the jobs incorrectly appear to run even when the job state is queued.</p>
BDM-22490	<p>When a dynamic mapping that runs on Spark uses an Update Strategy transformation, adding a column to the Hive target table schema causes the mapping to fail.</p>
BDM-22481	<p>When the Spark engine processes an input value of zero in a decimal port that is configured with equivalent precision and scale, the engine treats the value as data overflow and the return value is NULL.</p>
BDM-22282	<p>The Spark engine might take up to 30 minutes to run a mapping that contains a Python transformation if you pass a large number of ports to the Python transformation.</p>

Bug	Description
BDM-22260	Cannot get Spark monitoring statistics for a mapping run that uses any of the following connections: Google BigQuery, Google Cloud Storage, Google Cloud Spanner, and Google Analytics.
BDM-2222	The Spark engine does not run the footer row command configured for a flat file target. (459942)
BDM-2141	Mapping with a Hive source and target that uses an ABS function with an IIF function fails in the Hadoop environment. (424789)
BDM-2137	Mapping in the Hadoop environment fails when it contains a Hive source and a filter condition that uses the default table name prefixed to the column name. Workaround: Edit the filter condition to remove the table name prefixed to the column name and run the mapping again. (422627)
BDM-2136	Mapping in the Hadoop environment fails because the Hadoop connection uses 128 characters in its name. (421834)
BDM-20856	When you import a cluster workflow, the import wizard does not include the option to choose the non-native connection that was associated with the Create Cluster task. Workaround: After you import the workflow, manually assign a Databricks or Hadoop connection to the Create Cluster task.
BDM-20697	If you use the numberOfErrorRows system-defined mapping output in a mapping that runs on the Spark engine, the engine returns an incorrect value for the mapping output.
BDM-18140	A mapping that reads a large number of reference tables may take longer than expected to run on the Spark engine. The issue is observed when the mapping includes transformations that collectively read 140 reference tables. Workaround: Run the mapping on the Blaze engine.
BDM-17485	Mapping or mapplet that has the same name or any non-reusable transformation with the same name as mapplet or mapping fails to import into the Model repository.
BDM-17174	When memory usage reaches the maximum container size, YARN kills the container. Memory usage on the OOP Container Manager reaches the maximum container size if the following conditions are true: - Concurrent jobs run for longer than two days. - The Blaze engine has not reached the idle timeout limit or the sunset time.
BDM-16521	On Oracle and DB2, when a Lookup transformation contains a Text data type column and you import the mapping into the Developer tool, the Text datatype is mapped to the Clob datatype and the mapping fails with an error. Similarly, the Binary data type gets imported into the Model repository as a Blob data type and the mapping fails. Workaround: Edit the column type in the Lookup transformation to run the mapping.

Big Data Streaming Known Limitations

This following table describes known limitations:

Bug	Description
IIS-2781	A mapping with Hive target fails after 24 hours on a secure cluster.
IIS-2761	If a mapping contains multiple lookups on the same flat file data object, you cannot view the Spark execution plan in the Developer tool or Administrator tool. Workaround: Create a physical data object on the same physical source for each Lookup transformation in the pipeline.
IIS-2573	When you run a dynamic mapping and you connect the Window port from a Window transformation to an Aggregator transformation, the mapping validation fails.
IIS-2329	When you remove or update the AWS credentials file with an invalid credential, streaming mappings that use Amazon Kinesis connections hang.

Developer Tool Known Limitations

The following table describes known limitations:

Bug	Description
PLAT-17627	When you run multiple concurrent mappings with Lookup and Sorter transformations in the native environment, the mappings might hang.
PLAT-16416	When you import a mapping from PowerCenter that contains a maplet and an Expression transformation, the following error sometimes appear in the log: <code>The transformation instance cannot be connected to the active transformations.</code> Workaround: Ignore the error and import the mapping.
OCON-990	The Developer tool imports the DECFLOAT data type of IBM DB2 for z/OS tables as Char(0). (455216)
OCON-9806	If the ODBC data source contains a special character in a column name, mappings that read data from or write data to Microsoft SQL Server databases fail.
BDM-19301	The Incremental Deployment Wizard analyzes the impact to run-time objects one at a time, and it cannot perform an analysis in bulk. Analysis is slow for applications with a large number of objects.

Enterprise Data Catalog Known Limitations

The following table describes known limitations:

Bug	Description
PLAT-17013	The value frequency pane in the Assets Details view does not display minimum and maximum values for columns with Date data type.
EIC-8187	When you configure a user with read and write permission for one resource in the Catalog Administrator, Enterprise Data Catalog displays the total resource count instead of a count of one in the search page for the user.
EIC-8171	When you delete an Informatica Axon resource, Enterprise Data Catalog does not delete the custom attribute assigned to the asset from the catalog.
EIC-8127	When you search for an asset name that contains an underscore (_), the search results page does not return any result.
EIC-7934	Enterprise Data Catalog and Catalog Administrator display junk characters in the search results and resource properties section after you change the server locale to Japanese.
EIC-7741	The Column profile option is incorrectly enabled for unstructured data sources. Workaround: Use Data domain discovery option or Column profile and data domain discovery option for unstructured data sources.
EIC-7622	Custom metadata resource does not extract connection parameters for Business Intelligence metadata sources.
EIC-7583	The relationships REST API returns different levelCount details for the same query.
EIC-7406	After you fix the incorrect details in a reusable configuration used for a resource, the changes do not take effect on the resource.
EIC-7405	You cannot select a negative value using the slider filter for an integer custom attribute. Workaround: Configure the custom attribute to allow filter based on range. Draft Comment: Sabha S: Added in 10.2. Target release = 10.2.1. Tech reviewer: Ainn Rahathul
EIC-6927	A run-time error appears for a data domain or composite data domain if the following conditions are true:1. You open the Catalog Administrator in Internet Explorer version 11.2. You create and save a data domain or composite data domain. This issue occurs only for data domains and composite data domains that have a non-EN character in their names.
EIC-6084	The <code>ERROR: java.lang.IndexOutOfBoundsException</code> error message appears when you run the <code>infacmd collectAppLogs</code> command to collect log files of YARN applications. The error occurs when you create and enable the Catalog Service on an existing Hortonworks cluster that is enabled for WANdisco Fusion.
EIC-4015	When you run an Avro file with fixed data type, the core scanner displays the precision as undefined and the profile run fails with a null pointer exception.
EIC-3848	Profiling scanner does not run on Avro and Parquet files if the files are located in OneDrive, SFTP, or remote Windows location.
EIC-3288	When a column in a data source does not contain any data, the Similarity Discovery system resource does not discover columns based on column name.

Bug	Description
EIC-1837	Relationship diagram does not show synonyms.
EIC-1835	In the Asset Details view of a synonym, the Contains panel displays source qualifier and mapping assets along with columns instead of just columns.
EIC-15881	You cannot configure multi-realm Kerberos authentication for Enterprise Data Catalog.
EIC-15822	The <code>infacmd collectAppLogs</code> command, when used with the <code>-ad</code> (Application Start Date) option, fails with the <code>java.lang.NullPointerException</code> error. The error occurs when you run the command on an existing cluster enabled for Kerberos and WANdisco Fusion.
EIC-15755	The Data Domain Propagation resource fails with the following error: "ArrayIndexOutOfBoundsException". This issue applies to an IBM DB2 for z/OS resource that has Similarity Profile enabled.
EIC-15733	Enterprise Data catalog does not import all the multi-valued custom attribute values of the reference data type. This issue occurs if you import an Axon glossary into the catalog.
EIC-15696	The <code>infacmd collectAppLogs</code> command fails with the <code>collectAppLogs] failed with error [Java heap space]</code> error message.
EIC-15657	In Enterprise Data Catalog, you can create a Universal Connectivity Framework (UCF) resource using REST API. However, you cannot edit the resource. Workaround: Contact Global Customer Support for a UCF license.
EIC-15650	UCF resource model registration fails for SAP HANA resource after Catalog Service upgrades successfully.
EIC-15601	The Overview tab displays custom attributes even after you remove the custom attributes from the Application Configuration page.
EIC-15597	Connection assignment for the Amazon S3 metadata resource fails in Enterprise Data Catalog, if you override the Folder Path and File Name data objects of the metadata resource. Workaround: Do not override the Folder Path and File Name data objects of the Amazon S3 metadata resource.
EIC-15554	You cannot export or import the inferred and auto-accepted data domains.
EIC-15528	When you select more than one data domain and you curate it, only one data domain is curated and the other data domains are ignored. You cannot infer data domains if the data domain name contains special characters.
EIC-15510	Enterprise Data Catalog does not display the lineage correctly for a Hadoop Multi-Part file directory source and a relational target. Workaround: Use the Lineage bar, to manually expand multiple levels of lineage.
EIC-15338	In the Catalog Administrator, when you apply a custom attribute value to all assets in the resource, the following issues occur in the Notification page: <ul style="list-style-type: none"> - The Notification page displays duplicate notification messages. - The Notification page displays the notification message before the asset changes reflect in the Catalog. - The notification message displays the object ID of an asset instead of the asset name.
EIC-14949	When you run a File System resource with SFTP file protocol and First N rows sampling option, the profiling scanner scans all the rows in the resource.

Bug	Description
EIC-1486	The Transformation Logic view displays inaccurate mapping names for the Informatica Cloud Services transformations.
EIC-14792	Sometimes, Enterprise Data Catalog displays incorrect search results when you perform a wild card search.
EIC-14761	Sometimes, Enterprise Data Catalog displays an incorrect notification message in the Change Summary for a synonym asset.
EIC-14746	Enterprise Data Catalog does not send a notification message when a user associates or removes a business title from a stored procedure. This issue occurs if you have followed the enrichment changes of the stored procedure asset.
EIC-14693	Sometimes, you cannot save the newer login credentials for Enterprise Data Catalog Plugin if you have saved the credentials previously.
EIC-14668	In the Catalog Administrator, when you propagate the stakeholder value to all assets in the resource while you create a resource, the Relationship tab displays incorrect information for an asset.
EIC-14600	For custom metadata resources, you cannot specify multiple values for the multi-valued attributes that are configured in the Enterprise Data Catalog model.
EIC-14590	In the Catalog Administrator, the Manage menu does not display the Reusable Configuration option. This issue occurs only if you use Microsoft Internet Explorer version 11.1446.
EIC-14452	In the Enterprise Data Catalog, a user is unable to provide a review, ask a question, or answer a question about an asset if the username contains Chinese characters.
EIC-14429	Sometimes, Enterprise Data Catalog displays a previously assigned business title even after you update the business title for an asset.
EIC-14406	Enterprise Data Catalog does not display the path of an asset in the search results page. This issue occurs if you search for a column asset with keywords "column contains" followed by the asset name in the search box.
EIC-14374	Enterprise Data Catalog displays an incorrect notification message after you delete a description of an asset.
EIC-14183	The Resource Name filter option in the Assets Followed page does not work as expected.
EIC-14140	Sometimes, if you follow the collaboration changes of an asset, Enterprise Data Catalog does not send you a notification message.
EIC-14127	Enterprise Data Catalog stops unexpectedly when you type a percentage sign (%) after the asset name in the Asset Name field of the Notification page.
EIC-13718	The filters in the Relationship tab do not work as expected.
EIC-13708	Enterprise Data Catalog does not extract metadata or lineage information from an Informatica Intelligent Cloud Services (IICS) metadata source that includes an invalid data mapping with a Google BigQuery source. The invalid mapping between Google BigQuery and IICS is caused by IICS not supporting the Google BigQuery numeric data type.
EIC-13614	Enterprise Data Catalog displays an incorrect notification message after you delete a question.

Bug	Description
EIC-13577	After you assign a data owner to a resource and an asset, there is a user role and privileges conflict between the data owners.
EIC-13418	You cannot cancel a resource purge operation.
EIC-13353	In the Catalog, the average ratings count is incorrect when multiple users provide ratings and review at the same time.
EIC-13305	When you run a profile on a Salesforce resource, sometimes, the profiling fails and an error appears for a few tables in the resource.
EIC-13285 and EIC-13181	Enterprise Data Catalog does not display lineage for a view that includes advanced Google BigQuery functions and syntaxes such as ARRAY, UNNEST, CODE_POINTS_TO_BYTES, RANK, GROUP EACH BY and CODE_POINTS_TO_STRING.
EIC-13271	You cannot view the lineage for a view created in a Google BigQuery table if the following conditions are true: <ol style="list-style-type: none"> 1. You create a table in a Google BigQuery project, and create a resource for the table. 2. You create a view for the table in another Google Big Query project, and create a resource for the view. 3. You run both the resources. 4. Generate the lineage of the view.
EIC-13169	Lineage is not displayed in Enterprise Data Catalog for an Informatica Platform resource if the resource is mapped to a source or target file that has special characters such as ?, and * in the file name.
EIC-13164	The Notification page does not display a notification message when a CSV file is added to an asset that you follow.
EIC-13142	You cannot assign stakeholders to an asset if the asset name contains Chinese characters or UTF-8 characters.
EIC-13114	The lineage diagram of a synonym asset hangs unexpectedly if you drill down to display the lineage at the column level.
EIC-13002	Sometimes, after you remove an associated business title from an asset, the search results page displays the asset with the associated business title.
EIC-13002	Deleting an associated business glossary from a business glossary resource does not remove the associated business glossary term from the search results in Enterprise Data Catalog. This issue occurs for Business Glossary and Axon resource types.
EIC-12990	Enterprise Data Catalog does not send a notification message if you make changes to an asset that has special characters in its name.
EIC-12985	The Catalog Administrator and Enterprise Data Catalog stop unexpectedly in Microsoft Internet Explorer version 11.1446. This issue occurs when you create a custom attribute in Catalog Administrator and assign a custom attribute value to an asset in Enterprise Data Catalog.
EIC-12687	In the Library workspace, the resource filter section displays the Change Notification Resource option in the Resource Type list even if you have not selected the Enable Email Notifications for Asset Changes option in the Catalog Service properties of the Informatica Administrator.
EIC-12423	You cannot reject an inferred business term if the business term was previously associated with an asset more than two times.

Bug	Description
EIC-12413	When a business term is automatically assigned as a business title to a column that has an existing business title, the search results page displays both the business terms that are associated with the column. For example, the business terms might appear in the search results page in the following way: <BUSINESSTERM1>\$\$INFA_SEPARATOR\$\$<BUSINESSTERM2>
EIC-11876	When you create a single value or multi-value custom attribute of reference data type with the GET "http://<hostname>:<port>/access/2/catalog/models/referenceAttributes REST API, update the value in Enterprise Data Catalog, and refresh the page, the single value attribute disappears and the multi-value attribute displays all possible values.
EIC-11572	When you search for a few data sources and reports that are extracted to the catalog using a Workday resource type, Enterprise Data Catalog does not display the assets in search results. Workaround: Enclose the search string in double quotes (" ") when you perform a search in the Enterprise Data Catalog.

Enterprise Data Lake Known Limitations

The following table describes known limitations:

Bug	Description
IDL-4089	When you unshare a visualization Notebook, the asset details page shows the Zeppelin welcome page.
IDL-3615	When you share a visualization Notebook, the asset details page shows the Zeppelin welcome page.
IDL-3189	When you import data or preview data from Azure SQL Data Warehouse, the operation fails.

Profiles and Scorecards Known Limitations

The following table describes known limitations:

Bug	Description
IDE-3693	In the Analyst tool, drilldown on a column with Date data type fails when the following conditions are true: <ol style="list-style-type: none"> 1. You choose Oracle JDBC connection for the profiling warehouse. 2. You choose a JDBC connection to connect to a source to import a table. 3. You create and run a column profile on the table.
IDE-3178	When you enable operating system profile and drilldown on the column profiles or scorecard results, the drilldown sometimes fail.
IDE-3163	In Informatica Analyst, you cannot export profiles and scorecards if the pop-up blocker is enabled for the browser.
IDE-3155	The detail view for a column with data type Decimal takes a long time to appear if the profiling warehouse uses Microsoft SQL Server database.

Bug	Description
IDE-3114	In the Analyst tool, you need to check in a profile after every edit to run a profile but you can edit and run a scorecard multiple times before you check in the scorecard if Team-based development is enabled.
IDE-3089	An error appears after you perform a query on the MRX_PROFILE_RULES view if the Model repository uses Microsoft SQL Server and Microsoft Azure SQL database.
IDE-3086	Create, open, and run profiles or scorecards takes a long time if the Model repository uses Microsoft Azure SQL database.
IDE-3075	In the Developer tool, when you run a column profile on a JSON data object, some fields might not appear in the profile definition or profile results.
IDE-3063	An exception appears when you run a column profile on an Avro or Parquet data object in Amazon S3 bucket.
IDE-3061	An error appears when you create a column profile on an Avro data object that contains an array data type.
IDE-3057	When you run a column profile on a JSON data object, the special characters !, @, #, \$, %, ^, &, and * in the column field names are replaced with underscores in the profile definition and profile results.
IDE-2769	When you run a SQL query that uses the MRX_SC_RULE_METRICS view, the query results appear after an indefinite period of time.
IDE-2695	In the Developer tool, when you create a column profile on a logical data object and add a filter to the profile, the filter preview does not appear.

Reference Data Known Limitations

The following table describes known limitations:

Bug	Description
IDQ-6656	You cannot import an XML object to the Model repository when the following conditions are true: <ul style="list-style-type: none"> - The XML object includes a reference table. - The Model repository connects to an Oracle reference data database, and the reference table in the XML object originates in a Microsoft SQL Server reference data database.

Rule Specifications Known Limitations

The following table describes known limitations:

Bug	Description
IDQ-6355	The Analyst tool fails to invalidate a rule specification when you configure a case conversion function in a rule statement for a non-string data type.
IDQ-5317	The Analyst tool generates an error if you try to generate a mapplet from a rule specification after you perform the following steps: 1. You copy a rule statement from a parent rule set to a child rule set. 2. You delete the child rule set. Workaround: Save the current rule specification after you perform each copy operation.
IDQ-5316	The Analyst tool generates an error if you try to save a rule specification or generate a mapplet from a rule specification in the following case: - You copy a rule set from more than one other rule specification into the current rule specification. Workaround: Save the current rule specification after you perform each copy operation.

Third-Party Known Limitations

The following table describes known limitations:

Bug	Description
OCON-9943	If you configure Sqoop to import time data from a Netezza database, the mapping fails. Apache ticket reference number: SQOOP-2978
OCON-8786	If you configure Sqoop to export data of the Clob or DBClob data type to IBM DB2 z/OS targets, the mapping fails.
OCON-8561	If you configure Sqoop to export data of the Money data type to Microsoft SQL Server targets, the mapping fails.
OCON-8387	If you configure TDCH and Sqoop and run a mapping on the Blaze or Spark engine to export data of the Time data type, only milliseconds are written to the target. The nanosecond part is truncated. Cloudera ticket reference number: 124306
OCON-8332	If you configure Sqoop to export data of the Clob or DBClob data type to IBM DB2 targets, the mapping fails.
OCON-7974	If you configure Sqoop and a column name contains spaces, the mapping fails. Apache ticket reference number: SQOOP-2737
OCON-7620	If you import data from an IBM DB2 source through Sqoop and the table name contains mixed case characters, the mapping fails. Sqoop JIRA issue number: SQOOP-3211

Bug	Description
OCON-7505	<p>Sqoop mappings that read byte or varbyte data from a Teradata source and write it to a Teradata target fail on the Blaze engine. This issue occurs if you use Cloudera Connector Powered by Teradata.</p> <p>Cloudera ticket reference number: 124305</p>
OCON-7504	<p>When you use Sqoop to read data of the Timestamp data type from a Teradata source and write it to a Teradata target, only milliseconds are written to the target. This issue occurs if you run the Teradata Parallel Transporter mapping on a Cloudera cluster and on the Blaze engine.</p> <p>Cloudera ticket reference number: 124302</p>
OCON-7503	<p>When you use Sqoop to read time data from a Teradata source and write it to a Teradata target, the fractional seconds get corrupted. This issue occurs if you use Cloudera Connector Powered by Teradata or Hortonworks Connector for Teradata, and run the mapping on the Blaze engine.</p> <p>Cloudera ticket reference number: 124306</p>
OCON-7459	<p>When you export data to an IBM DB2 target through Sqoop, the mapping fails if all of the following conditions are true:</p> <ul style="list-style-type: none"> - You create or replace the IBM DB2 target table at run time. - The IBM DB2 target table name or column names contain mixed case characters. - You run the mapping on a Cloudera 5u8 cluster. <p>Apache ticket reference number: SQOOP-3212</p>
OCON-7431	<p>When you read time data from a Teradata source and write it to a Teradata target, the fractional seconds get corrupted. This issue occurs if you run the Teradata Parallel Transporter mapping on a Hortonworks cluster and on the Blaze engine.</p> <p>Cloudera ticket reference number: 124302</p>
OCON-7219	<p>When you run a Sqoop mapping on the Blaze engine to export Teradata float data, the data is truncated after the decimal point.</p> <p>Cloudera support ticket number: 113716</p>
OCON-7214	<p>Sqoop mappings fail on the Blaze engine if you use a custom query with the Order By clause to import data.</p> <p>Sqoop JIRA issue number: SQOOP-3064</p>
OCON-7213	<p>The Sqoop program does not honor the --num-mappers argument and -m argument when you export data and run the mapping on the Blaze or Spark engine.</p> <p>Sqoop JIRA issue number: SQOOP-2837</p>
OCON-7211	<p>When you run a Sqoop mapping to import data from or export data to Microsoft SQL Server databases that are hosted on Azure, the mapping fails.</p> <p>Sqoop JIRA issue number: SQOOP-2349</p>
OCON-2847	<p>Loading a Microsoft SQL Server resource fails when TLS encryption is enabled for the source database and the Metadata Manager repository is a Microsoft SQL Server database with TLS encryption enabled. (452471)</p> <p>Data Direct case number: 00343832</p>

Bug	Description
OCON-1100	<p>When you export data to an IBM DB2 z/OS database through Sqoop and do not configure the batch argument, the mapping fails.</p> <p>Workaround: Configure the batch argument in the mapping and run the mapping again. (459671)</p> <p>Apache ticket reference number: SQOOP-2980</p>
IIS-2719	<p>If the mapping contains a passive Lookup transformation that is configured with an inequality lookup condition, the target receives the data later.</p>
IIS-2484	<p>When you write to a HBase target in an Amazon EMR cluster and the spark.hadoop.validateOutputSpecs property is set to true, the streaming mapping fails with an exception.</p> <p>Workaround: Set the spark.hadoop.validateOutputSpecs property to false.</p> <p>HBASE JIRA issue number: 20295</p>
BDM-9585	<p>Mappings fail on the Spark engine when you configure an SQL override to access a Hive view.</p> <p>Apache Spark ticket reference number: SPARK-21154.</p>
BDM-23420	<p>A mapping with a Sorter transformation that contains datetime data in its default value changes the data when you run the mapping on a MapR cluster.</p> <p>MapR ticket reference number: 00072094</p>
BDM-23104	<p>The Spark engine cannot write data to a bucketed Hive target if the Hadoop distribution is MapR.</p> <p>MapR case number: 00074338</p>
BDM-22675	<p>Incorrect source statistics appear when the Spark engine processes a Union transformation.</p> <p>Spark ticket reference number: SPARK-26683</p>
BDM-22318	<p>If a parameter default value has a newline character, the Spark engine creates additional rows and can shift column values to other columns.</p> <p>Workaround: To resolve the issue with newline characters, use a SequenceFile format instead of a TextFile format for the target table.</p>
BDM-17470	<p>In an Azure HDInsight environment, if you enable Hive merge in an Update Strategy transformation and if Hive is enabled to execute vectorized queries, inserting data into specific columns fails.</p> <p>Workaround: In the hive-site.xml on the cluster, set the hive.vectorized.execution.enabled property to false.</p> <p>Apache Hive ticket reference number: HIVE-14076</p>
BDM-17204	<p>If the impersonation user tries to run a mapping with HDFS source and does not have the DECRYPT_EEK privilege to read the file, the log file shows an incorrect error message.</p> <p>HADOOP-12604</p>
BDM-17020	<p>When you run a mapping that uses a schema in an Avro file, the Spark engine adds a NULL data type to the primitive data types in the schema.</p>
BDM-14438	<p>If a Parquet file has an array of the format [map (struct, any)], then the mapping on the Spark engine fails with a spark.sql.AnalysisException exception.</p> <p>SPARK-22474</p>

Bug	Description
BDM-14422	The mapping fails with an error on Spark engine due to duplicate columns in the Hive table. SPARK-23519
BDM-14410	The mapping fails because the Spark engine cannot read from an empty ORC Hive source. SPARK-19809
BDM-10570	The Spark job fails with out of memory errors when a mapping that converts relational data to hierarchical data contains more than three Aggregator and Joiner transformations. Workaround: To convert relational data to a hierarchical data of more than four levels, develop more than one mapping to stage the intermediate data. For example, develop a mapping that converts relational data to a hierarchical data up to three levels. Use the hierarchical data in another mapping to generate a hierarchical data of four levels. SPARK-22207
BDM-10455	Inserts into a bucketed table can sometimes fail when you use Hive on Tez as the execution engine. The issue is more probable if the table is a Hive ACID table and a delete operation is performed before the inserts. Apache ticket reference number: TEZ-3814

Transformations Known Limitations

The following table describes known limitations:

Bug	Description
IDQ-6914	If you specify a reference table name as a parameter in a transformation, the transformation disregards the reference table name when you resolve the mapping parameters in the mapping canvas.
IDQ-6664	A mapping that contains a Labeler transformation or a Parser transformation fails validation when the following conditions are true: <ul style="list-style-type: none"> - You configure the transformation to read a probabilistic model from a content set. - You replace the probabilistic model in the content set with a probabilistic model of the same name from another content set. When the conditions are true, the probabilistic model no longer appears in the transformation strategy.
IDQ-6662	A mapping that contains a Classifier transformation fails validation when the following conditions are true: <ul style="list-style-type: none"> - You configure the transformation to read a classifier model from a content set. - You replace the classifier model in the content set with a classifier model of the same name from another content set. When the conditions are true, the classifier model no longer appears in the transformation strategy.

Workflows Known Limitations

The following table describes known limitations:

Bug	Description
MWF-1430	When you assign the system parameter sys:ApplicationName or sys:MappingName to a Mapping task input in a workflow, the Mapping task input parameter does not read the system parameter values. Instead, the Mapping task input parameter reads the default value of the mapping parameter at run time.
MWF-1340	<p>If a workflow contains parallel Mapping tasks and you cancel the workflow while the Mapping tasks are running, any mapping that has yet to start proceeds to run.</p> <p>The issue arises when the following conditions are true:</p> <ul style="list-style-type: none">- The number of mappings that the Mapping tasks specify is greater than the Maximum Worker Threads value on the Data Integration Service.- You cancel the workflow while the mappings that the Maximum Worker Threads value specifies are running and other mappings have yet to start. <p>Workaround: Increase the Maximum Worker Threads value.</p>

Emergency Bug Fixes Merged into 10.2.2 Service Pack 1

Informatica merged Emergency Bug Fixes (EBFs) into version 10.2.2 Service Pack 1. These EBFs provide fixes for issues that were found in previous releases.

For a list of EBFs that were merged into version 10.2.2 Service Pack 1, see the following Informatica Knowledge Base article: <https://kb.informatica.com/faq/7/Pages/23/575963.aspx>.

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.