



Contents

Technical Preview.	2
Technical Preview for Big Data Management.	2
Technical Preview for Big Data Streaming.	3
Technical Preview for Enterprise Data Catalog.	3
Technical Preview for PowerExchange Adapters.	4
Installation and Upgrade.	5
Informatica Upgrade Paths.	5
Support of Informatica Big Data Products in 10.2.1.	6
Big Data Hadoop Distribution Support.	6
Hive Run-Time Engine	8
Upgrading to New Configuration.	9
10.2.1 Fixed Limitations and Closed Enhancements.	9
Administrator Tool Fixed Limitations (10.2.1).	9
Analyst Tool Fixed Limitations (10.2.1).	10
Application Service Fixed Limitations (10.2.1).	10
Big Data Management Fixed Limitations (10.2.1).	10
Big Data Streaming Fixed Limitations (10.2.1).	14
Command Line Programs Fixed Limitations (10.2.1).	15
Data Transformation Fixed Limitations (10.2.1).	15
Domain Fixed Limitations (10.2.1).	16
Enterprise Data Catalog Fixed Limitations (10.2.1).	16
Enterprise Data Lake Fixed Limitations (10.2.1).	17
Exception Management Fixed Limitations (10.2.1).	18
Mappings and Workflows Fixed Limitations.	18
Profiles and Scorecards Fixed Limitations (10.2.1).	20
Security Fixed Limitations (10.2.1).	21
Third-Party Fixed Limitations (10.2.1).	21
Transformation Fixed Limitations (10.2.1).	22
Transformation Language Functions Fixed Limitations (10.2.1).	22
10.2.1 Known Limitations.	22

Administrator Tool Known Limitations (10.2.1).	23
Analyst Tool Known Limitations (10.2.1).	23
Application Service Known Limitations (10.2.1).	23
Big Data Management Known Limitations (10.2.1).	25
Big Data Streaming Known Limitations (10.2.1).	31
Domain Known Limitations (10.2.1).	32
Developer Tool Known Limitations (10.2.1).	32
Enterprise Data Catalog Known Limitations (10.2.1).	34
Enterprise Data Lake Known Limitations (10.2.1).	37
Parameters Known Limitations (10.2.1).	37
Profiles and Scorecards Known Limitations (10.2.1).	37
Reference Data Known Limitations (10.2.1).	38
Rule Specification Known Limitations (10.2.1).	39
Security Known Limitations (10.2.1).	39
Third-Party Known Limitations (10.2.1).	40
Transformations Known Limitations (10.2.1).	42
Workflows Known Limitations (10.2.1).	43
Informatica Global Customer Support.	43

This document contains important information about restricted functionality, known limitations, and bug fixes for Informatica 10.2.1.

Technical Preview

Technical Preview for Big Data Management

Big Data Management version 10.2.1 includes functionality that is available for technical preview. Technical preview functionality is supported 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 GA 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.

Intelligent Structure Discovery

Effective in version 10.2.1, the intelligent structure model is available for technical preview in Amazon S3 data objects, complex file data objects, and Microsoft Azure Blob data objects.

You can read files that are parsed with an intelligent structure model when you run a mapping on the Spark engine. When you add the data object to a mapping that runs on the Spark engine, you can process any input type that the model can parse.

Python Transformation

Effective in version 10.2.1, the Python transformation is available for technical preview. Use the Python transformation to execute Python code in a mapping that runs on the Spark engine.

You can use a Python transformation to implement a machine model on the data that you pass through the transformation. For example, use the Python transformation to write Python code that loads a pre-trained model. You can use the pre-trained model to classify input data or create predictions.

Technical Preview for Big Data Streaming

Big Data Streaming version 10.2.1 includes functionality that is available for technical preview. Technical preview functionality is supported 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 GA 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.

Python Transformation

Effective in version 10.2.1, the Python transformation is available for technical preview. Use the Python transformation to execute Python code in a mapping that runs on the Spark engine.

You can use a Python transformation to implement a machine model on the data that you pass through the transformation. For example, use the Python transformation to write Python code that loads a pre-trained model. You can use the pre-trained model to classify input data or create predictions.

Complex Data Type in Java Transformation

Effective in version 10.2.1, when you use the Java transformation, you can use complex data types to process hierarchical data. The complex data type support in the Java transformation is available for technical preview.

Map Data Type

Effective in version 10.2.1, you can use map data type to generate and process map data on the Spark engine. The map data type support is available for technical preview.

Technical Preview for Enterprise Data Catalog

Enterprise Data Catalog version 10.2.1 includes functionality that is available for technical preview. Technical preview functionality is supported 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 GA 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.

Import from ServiceNow

Effective in version 10.2.1, Catalog Administrator now connects to ServiceNow to import connections and extract the configuration metadata into the catalog. The Import from ServiceNow feature is available for technical preview.

Database Script Resource Type

Effective in version 10.2.1, you can create a Database Script resource to extract lineage information from database scripts. The Database Script resource type is available for technical preview.

Technical Preview for PowerExchange Adapters

PowerExchange Adapters version 10.2.1 includes functionality that is available for technical preview. Technical preview functionality is supported 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 GA release for production use, but might choose not to in accordance with changing market or technical circumstances. For more information, contact Informatica Global Customer Support.

PowerExchange for Amazon S3

Effective in version 10.2.1, PowerExchange for Amazon S3 includes the following technical preview functionality:

- You can read files that are parsed with an intelligent structure model when you run a mapping on the Spark engine. When you add the data object to a mapping that runs on the Spark engine, you can process any input type that the model can parse.
- You can use Amazon S3 sources as dynamic sources in a mapping.

For more information, see the *Informatica PowerExchange for Amazon S3 10.2.1 User Guide*.

PowerExchange for HDFS

Effective in version 10.2.1, you can incorporate an intelligent structure model in a complex file data object. When you add the data object to a mapping that runs on the Spark engine, you can process any input type that the model can parse.

For more information, see the *Informatica PowerExchange for HDFS 10.2.1 User Guide*.

PowerExchange for Microsoft Azure Blob Storage

Effective in version 10.2.1, PowerExchange for Microsoft Azure Blob Storage includes the following technical preview functionality:

- You can run mappings on the Spark engine.
- You can read and write .csv, Avro, and Parquet files when you run a mapping on the Spark engine and in the native environment.
- You can read and write JSON when you run a mapping on the Spark engine.
- You can read files that are parsed by an intelligent structure model when you run a mapping on the Spark engine.
- You can read a directory when you run a mapping on the Spark engine.
- You can generate or skip header rows when you run a mapping in the native environment. On the Spark engine, the header row is created by default.
- You can append an existing blob. The append operation is applicable to only to the append blob and in the native environment.

- You can override the blob or container name. In the Blob Container Override field, specify the container name or sub-folders in the root container with the absolute path.
- You can read and write .csv files compressed in the gzip format.

For more information, see the *Informatica PowerExchange for Microsoft Azure Blob Storage 10.2.1 User Guide*.

Installation and Upgrade

Informatica Upgrade Paths

You can directly upgrade to 10.2.1 from Informatica 10.0, 10.1, 10.1.1, and 10.2.

If the product version that is currently installed cannot be upgraded to Informatica 10.2.1, you must first upgrade to a supported version. To determine the Informatica product version that is currently installed, click **Help > About Informatica Administrator** in the Informatica Administrator header area.

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

The following table describes the Informatica product versions from which you can upgrade:

Informatica Version	Upgrade Path	Comments
9.0.1	9.1.0 -> 9.6.1 -> 10.2 ->10.2.1	After you upgrade to version 9.1.0, you must upgrade to version 9.6.1, then upgrade to version 10.2, and then upgrade to version 10.2.1.
9.1.0	9.6.1 -> 10.2 ->10.2.1	After you upgrade to version 9.6.1, you must upgrade to version 10.2, and then upgrade to version 10.2.1.
9.5.0	9.6.1 -> 10.2 ->10.2.1	After you upgrade to version 9.6.1, you must upgrade to version 10.2, and then upgrade to version 10.2.1.
9.5.1	9.6.1 -> 10.2 ->10.2.1	After you upgrade to version 9.6.1, you must upgrade to version 10.2, and then upgrade to version 10.2.1.
9.6.0	9.6.1 -> 10.2 ->10.2.1	After you upgrade to version 9.6.1, you must upgrade to version 10.2, and then upgrade to version 10.2.1.
9.6.1	10.2 ->10.2.1	You must first upgrade to version 10.2 and then upgrade to version 10.2.1
10.0.0	10.2.1	You can directly upgrade to version 10.2.1.
10.1.0	10.2.1	You can directly upgrade to version 10.2.1.

Informatica Version	Upgrade Path	Comments
10.1.1	10.2.1	You can directly upgrade to version 10.2.1.
10.2.0	10.2.1	You can directly upgrade to version 10.2.1.

Important: Version 10.2 HotFix 1 was released after version 10.2.1. Informatica does not support the upgrade from 10.2.1 to 10.2 HotFix 1, nor does it support the upgrade from 10.2 HotFix 1 to 10.2.1. The upgrade process might succeed, but you might experience results that are not supported.

Support of Informatica Big Data Products in 10.2.1

Version 10.2.1 supports Informatica big data products, such as Big Data Management and Big Data Quality. It does not support traditional products such as PowerCenter and Informatica Data Quality. If your domain contains both big data and traditional products, you must split the domain before you perform the upgrade.

Big Data Hadoop Distribution Support

Informatica big data products support a variety of Hadoop distributions. In each release, Informatica adds, defers, and drops support for Hadoop distribution versions. Informatica might reinstate support for deferred versions in a future release.

The following table lists the supported Hadoop distribution versions for Informatica 10.2.1 big data products:

Product	EMR	HDI	CDH	HDP	MapR
Big Data Management	5.10, 5.14 ³	3.6.x	5.11 ¹ , 5.12 ¹ , 5.13, 5.14, 5.15	2.5, 2.6	6.x MEP 5.0.x ²
Big Data Streaming	5.10, 5.14 ³	3.6.x	5.11 ¹ , 5.12 ¹ , 5.13, 5.14, 5.15	2.5, 2.6	6.x MEP 4.0.x
Enterprise Data Catalog	N/A	3.6.x	5.13	2.6.x	N/A
Enterprise Data Lake	5.10	3.6.x	5.13	2.6.x	N/A

¹ Big Data Management and Big Data Streaming support for CDH 5.11 and 5.12 requires EBF-11719. See KB article [533310](#).

² Big Data Management support for MapR 6.x with MEP 5.0.x requires EBF-12085. See KB article [553273](#).

³ Big Data Management and Big Data Streaming support for Amazon EMR 5.14 requires EBF-12444. See KB article [560632](#).

Note: Informatica dropped support for IBM BigInsights.

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

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

Big Data Management Hadoop Distributions

The following table lists the supported Hadoop distribution versions and changes in Big Data Management 10.2.1:

Hadoop Distribution	Supported Distribution Versions	10.2.1 Changes
Amazon EMR	5.10, 5.14	Added support for version 5.10 and 5.14. Dropped support for version 5.8.
Azure HDInsight	3.6.x	Added support for version 3.6.x. Dropped support for 3.5x.
Cloudera CDH	5.11, 5.12, 5.13, 5.14, 5.15	Added support for versions 5.13, 5.14, 5.15.
Hortonworks HDP	2.5.x, 2.6.x	Added support for version 2.6.x. Dropped support for version 2.4.x.
MapR	6.x MEP 5.0.x	Added support for versions 6.x MEP 5.0.x. Dropped support for versions 5.2 MEP 2.0.x, 5.2.MEP 3.0.x.

Note: Informatica dropped support for IBM BigInsights.

Informatica big data products support a variety of Hadoop distributions. In each release, Informatica adds, defers, and drops support for Hadoop distribution versions. Informatica might reinstate support for deferred versions in a future release.

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>

Big Data Streaming Hadoop Distributions

The following table lists the supported Hadoop distribution versions and changes in Big Data Streaming 10.2.1:

Hadoop Distribution	Supported Distribution Versions	10.2.1 Changes
Amazon EMR	5.10, 5.14	Added support for 5.10, 5.14. Dropped support for version 5.4.
Azure HDInsight	3.6.x	Added support for version 3.6.x.
Cloudera CDH	5.11, 5.12, 5.13, 5.14, 5.15	Added support for versions 5.13, 5.14, 5.15.

Hadoop Distribution	Supported Distribution Versions	10.2.1 Changes
Hortonworks HDP	2.5.x, 2.6.x	Added support for version 2.6.x. Dropped support for version 2.4.x.
MapR	6.x MEP 4.0.x	Added support for versions 6.x MEP 4.0. Dropped support for versions 5.2 MEP 2.0.x, 5.2.MEP 3.0.x.

Informatica big data products support a variety of Hadoop distributions. In each release, Informatica adds, defers, and drops support for Hadoop distribution versions. Informatica might reinstate support for deferred versions in a future release.

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>.

Hive Run-Time Engine

Effective in version 10.2.1, the MapReduce mode of the Hive run-time engine is deprecated, and Informatica will drop support for it in a future release. The Tez mode remains supported.

Mapping

When you choose to run a mapping in the Hadoop environment, the Blaze and Spark run-time engines are selected by default.

Previously, the Hive run-time engine was also selected.

If you select Hive to run a mapping, the Data Integration Service will use Tez. You can use the Tez engine only on the following Hadoop distributions:

- Amazon EMR
- Azure HDInsight
- Hortonworks HDP

In a future release, when Informatica drops support for MapReduce, the Data Integration Service will ignore the Hive engine selection and run the mapping on Blaze or Spark.

Profiles

Effective in version 10.2.1, the Hive run-time engine is deprecated, and Informatica will drop support for it in a future release.

The Hive option appears as Hive (deprecated) in Informatica Analyst, Informatica Developer, and Catalog Administrator. You can still choose to run the profiles on the Hive engine. Informatica recommends that you choose the Hadoop option to run the profiles on the Blaze engine.

Upgrading to New Configuration

After you move from a Microsoft SQL server custom schema to an SQL Server database enabled with trusted connection, the test connection fails with the following error:

```
Login failed for user 'UserName'
```

(PLAT-8450, 460338)

10.2.1 Fixed Limitations and Closed Enhancements

Administrator Tool Fixed Limitations (10.2.1)

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

The following table describes fixed limitations:

Bug	Description
IDQ-5941	When you create a Content Management Service, the current Data Integration Services on the domain might not appear in the service creation wizard. The issue arise when the following conditions are true: <ul style="list-style-type: none">- The domain contained a node that a user deleted before you began to create the Content Management Service.- The node that the user deleted included a Data Integration Service.
PLAT-20123	Double login appears in the Administrator tool.
PLAT-18767	The Administrator tool fails to start correctly because of a duplicate 'org.apache.commons.logging.Log' file.
PLAT-17815	After you log in to the Administrator tool, the tool requires you to log in again when you access the Domain view on the Manage tab or the Monitor view.

The following table describes closed enhancement requests:

Bug	Description
PLAT-19331	The Scheduler Service now automatically adjusts to Daylight Savings Time.
PLAT-19268	You can now delete custom properties in the Administrator tool.

Analyst Tool Fixed Limitations (10.2.1)

The following table describes fixed limitations:

Bug	Description
ANT-1202	When you try to open a mapping specification, the Analyst tool displays the following error: An unknown error has occurred. Check the Analyst logs or contact your administrator.
SS-50	After you upgrade to 10.1.1 HF1, the Search service index file grows too large.

Application Service Fixed Limitations (10.2.1)

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

The following table describes fixed limitations:

Bug	Description
SS-45	The search result do not display objects in the expected order when the search relevancy is calculated based of inverse document frequency. When the frequency of a term is low in the search index, the inverse document frequency parameter value becomes higher.
PLAT-20403	DTM threads that process data on the Data Integration Service consume high CPU.
PLAT-19208	A Data Integration Service that runs on a grid fails when the home directory is a shared location between nodes on the grid.
MRS-1577	Creating a search index consumes maximum Java heap size resulting in high CPU usage for Model repository.
MRS-1486	An error appears when you run the infacmd RestoreDomain command.
MRS-1452	The Authentication failed because the authentication token is invalid error appears frequently in the Model Repository Service log. When this error appears, you cannot perform any task in the Model repository unless you restart the Model Repository Service.
MRS-1431	When you perform a search on the Model repository content, the maximum number of objects that are returned is based on the total number of Model repository objects and not the total number of mapping objects in the repository.
BDM-11629	If the infa_rpm.tar is deleted from HDFS, the Data Integration Service incorrectly determines that the infa_rpm.tar file already exists in HDFS when you run subsequent mappings.

Big Data Management Fixed Limitations (10.2.1)

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

The following table describes fixed limitations:

Bug	Description
PLAT-20370	When you run a mapping task in a workflow, the Big Data Management log collector reports that the mapping logs are not found.
PLAT-18717	The Administrator tool fails to display a cluster configuration when one of the *-site.xml configuration sets contains a dot character (.) in the name, excluding the required character for the file extension.
PLAT-14603	When you enable the Model Repository Service, the following message no longer appears in the JSF logs until you enable the debug logs: [COMPONENT_0005] Failed to load component 'PortalsProvider'.
OCON-13045	A Sqoop mapping that contains an Oracle source and a Hive target writes NULL values in Decimal columns as zeroes in the target. This issue occurs if you run the Sqoop mapping on the Spark engine.
OCON-11600	If you use Kerberos clusters, Sqoop mapping logs do not provide adequate information to troubleshoot errors.
IDQ-6183	When a mapping with a transformation that reads a reference table runs in the Hive engine, the transformation reads a null value in the reference table as "/N" and returns the string "/N" as a valid value.
BDM-8517	The infacmd ms RunMapping command does not return the job ID.
BDM-6754	When the Data Integration Service is configured to run with operating system profiles and you push the mapping to an HDInsight cluster with ADLS as storage, the mapping fails with the following error: Exception Class: [java.lang.RuntimeException] Exception Message: [java.io.IOException: No FileSystem for scheme: adl]. java.lang.RuntimeException: java.io.IOException: No FileSystem for scheme: adl
BDM-2181	The summary and detail statistics empty for mappings run on Tez. (452224)
BDM-17696	When you set the tracing level to verbose in the Lookup transformation, the Data Integration Service does not write additional details to the log.
BDM-17559	When the Hive engine calculates an expression that uses decimal data types, the return value of the expression might be incorrect if the expression contains a function that honors precision and scale when evaluating arguments such as TO_CHAR or CAST.
BDM-17146	A mapping fails on the Spark engine in an SSL-enabled cluster if the mapping contains a Hive table that uses an HBase storage handler.
BDM-17100	The output file for an AWS S3 target contains a duplicate header line when the File Merge option is selected.
BDM-16971	When you run a mapping with a logical data object or flat file sources and if the mapping contains a Sorter or Joiner transformation, an error appears and the Integration Service shuts down unexpectedly.
BDM-16960	When you run a data preview or mapping as an OS profile user under the supplementary group access list, a permission denied error occurs.

Bug	Description
BDM-16955	When you run a mapping in native mode and a Filter transformation is present after an output stream of Normalizer, the data coming out of other output streams do not have a filter and incomplete data appears in the target.
BDM-16952	When you use Lookup in a mapping, the Integration Service shuts down intermittently with an assert in lkupcmgr.cpp (line # 322).
BDM-16848	The Spark engine converts leap year dates incorrectly when the date appears in the format MM-DD-YYY. For example, the Spark engine converts 02-29-2016 to 03-01-2016.
BDM-16782	Web service requests that contain an Aggregator transformation or Rank transformation and that reuse a DTM instance return incorrect responses.
BDM-16750	The Spark engine fails to generate an execution plan to create a partitioned Hive table using a custom DDL query.
BDM-16582	When you run a mapping that uses a complex file reader on the Spark engine, a memory leak causes increasing CPU usage on the Data Integration Service.
BDM-15736	When you use the Run dialog option to run a mapping, you cannot choose a different Data Integration Service to run the mapping from.
BDM-15234	Mapping generation through SQL queries takes more than 20 minutes.
BDM-15120	The ObjectIECmdImpl.java reads the INFA_OIE_VM_ARG system property before the Xmx property.
BDM-14952	Concurrent mappings that are generated at the same millisecond have same log file names with .log and .log.1 extension.
BDM-14554	An HDFS mapping that runs in the native environment against a Hadoop cluster that uses Kerberos authentication fails when it runs for a long time.
BDM-14526	After upgrading to 10.1.1 HotFix 1, when the mapping contains multiple Lookup transformations with the Lookup cache enabled, incorrect output appears and the Integration Service sometimes shuts down unexpectedly.
BDM-14522	When you validate an constant operand to a NULL expression in the Aggregator transformation, the validation fails with an error.
BDM-14348	A mapping fails on the Spark engine if the mapping creates a table in a non-default schema.
BDM-14334	Unable to output the complete list of files extracted to local data nodes during installation.
BDM-14153	The memory used to store a long and complex Hive query for string concatenation is too high.
BDM-14141	After an upgrade, mappings with a parameterized HDFS connection fail with a connection validation error.
BDM-14131	Validating a mapping to run on the Hive or Spark engine does not resolve dynamic ports for an Update Strategy transformation.
BDM-14113	After you upgrade from version 9.6.1 HotFix 3 to 10.2, the Decimal data type truncates the trailing zeroes after the Decimal point.

Bug	Description
BDM-13669	Mappings running on the Blaze engine cause a cluster failure on a cluster with more than 120 nodes.
BDM-13570	The Blaze engine fails with exit code 134 when running a profile.
BDM-13387	When running column profiling on Blaze, the mapping fails with an error like: "Invalid mapping... <target name> has [<x>] partition(s), but it must have [<y>] partition(s) ... mismatch between sorter and source."
BDM-12884	You can set environment variable INFA_EXTERNAL_IMPERSONATION_PROCESS variable in the OSP-supported Informatica services where you want to use your custom impersonation process.
BDM-12369	Tez fails to connect to the Application Timeline Server when you run a job on a Hortonworks HDP cluster that is enabled with Kerberos and SSL.
BDM-12163	A mapping with an unconnected Lookup transformation fails to generate an execution plan when it runs on the Blaze engine.
BDM-12361	When you create a view for a mapping that has sources with an SQL override, the fields are misaligned.
BDM-11927	When you import a cluster configuration and connectivity between the cluster and Informatica clients stops, the error message is not informative enough.
BDM-11830	When you run jobs on a grid and the connection between the Data Integration Service and connection service fails, the jobs that are further submitted to the Data Integration Service do not get executed.
BDM-11692	When you run mapping with multiple Hive sources and use a join in a query override, the data preview on a mapping specification target fails with a Java heap space error message.
BDM-11631	Using the Run Preview function for a Data Processor transformation in a cluster configuration triggers reinstallation of RPM on the cluster.
BDM-11392	When a mapping runs on the Blaze engine, the mapping reads data from Hive sources incorrectly if the physical data object and Hive source ports are not in sync.
BDM-11375	Mappings that run on the Spark engine fail because the hadoop.rpc.protection property is not correctly imported into the cluster configuration. The following error message appears: "No common protection layer between client and server."
BDM-11172	A fetch execution plan running in the Hive environment hangs for a mapping containing a Data Processor transformation with a very large number of ports.
BDM-11111	You cannot run a mapping with Sequence Generator transformation concurrently with other mappings on Data Integration Service that runs on a grid.
BDM-11109	Sqoop mappings that read binary data from Teradata and write to HBase fail.
BDM-11106	When you run a mapping that processes more than 1 GB of data on an HDInsight cluster, the Blaze engine duplicates the number of rows in the target.
BDM-11081	The Spark and Hive engines do not delete temporary directories created to run mappings.

Bug	Description
BDM-11049	After the Data Integration Service restarts, a mapping that contains a Data Processor transformation might fail to run on the Hive engine the first time that it executes.
BDM-10924	A mapping with an SQL query defined on a Hive source and target table, or with an SQL query generated on the Hive source, fails with an error like: FAILED: ParseException line <line number> Failed to recognize predicate '<reserved keyword>'. Failed rule: 'identifier' in expression specification
BDM-10897	When you create or import type definition libraries using JSON objects with key names that have the dot operator (.), the following validation error occurs: Evaluation failed and was not completed. Check the Developer tool logs for details.
BDM-10878	When a mapping runs on the Blaze engine, the Hive HDFS reader improperly uses a staging directory in the source HDFS location.
BDM-10837	Mapping performance is slow when the Spark engine writes to a partitioned Hive table on Amazon S3.
BDM-10670	Memory leaks occur on the Data Integration Service machine when the Spark engine heap memory reaches 4 GB.
BDM-10566	Mappings fail when you use the impersonation user to read or write files within an HDFS encryption zone in a Hadoop cluster and operating system profiles are enabled for the Data Integration Service.
BDM-10438	The option to truncate external partitioned Hive targets is not supported in BDM 10.1.1.
BDM-10098	Mappings with a Normalizer transformation fail on the Spark engine if the value in the Occurs column is large.

Big Data Streaming Fixed Limitations (10.2.1)

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

The following table describes fixed limitations:

Bug	Description
IIS-1567	When you import complex data type definitions for a Streaming mapping, the Import Complex Data Type Definitions dialog box does not display XML in the file format list.
IIS-1152	When you create a write data object operation for a MapR Streams data object and try to view the advanced properties tab, a null pointer exception occurs.

Bug	Description
IIS-1121	When you run applications on a MapR cluster, failover does not happen and the following error occurs: ERROR StreamsListener fs/client/marlin/cc/listener/listenerimpl.cc:699 Thread: 26778 Seek called on unsubscribed partitions
IIS-1062	When you run two concurrent Streaming mappings that read from the same Kafka broker, where one mapping contains a Union transformation and one mapping contains a Joiner transformation, the mappings do not run and a warning appears in the logs as follows: WARN ConsumerCoordinator: Auto offset commit failed for group IIS-Kafka-DStream-Consumer-1499852337694: Commit cannot be completed since the group has already rebalanced and assigned the partitions to another member.

Command Line Programs Fixed Limitations (10.2.1)

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

The following table describes fixed limitations:

Bug	Description
PLAT-20525	Nodes might hang and jobs might fail due to high CPU usage when running jobs from the command line.
PLAT-18380	The node.log file contains excessive messages with message codes LGS_10071 and LGS_10070.
PLAT-14268	Running multiple infacmd command line programs at the same time consumes excessive memory.

The following table describes closed enhancements:

Bug	Description
PLAT-19545	A new field disablePaging is added to the LDAP option of the infacmd isp SetLDAPConnectivity command which allows you to disable paging.

Data Transformation Fixed Limitations (10.2.1)

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

The following table describes fixed limitations:

Bug	Description
B2BT-262	A Data Processor transformation might fail when it has a mapper with a preprocessor that contains the TransformByParser component.
B2BT-261	A Hierarchical to Relational Data Processor transformation for a Library might fail unexpectedly and produce the following unrelated error: There are [X] rows in group [X] that are not in the output xml. Verify that the data is sorted correctly if the input is sorted. Verify that a primary key in the parent row is not NULL.

Domain Fixed Limitations (10.2.1)

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

The following table describes fixed limitations:

Bug	Description
PLAT-20728	The Alert Service fails when the SMTP server is configured without authentication.
PLAT-18622	When you purge a domain that contains a large number of files that must be purged, the domain becomes unavailable due to a lack of heap memory.
PLAT-11157	When the domain fails to respond to a connection initialization message within the timeout period, the Data Transport Framework incorrectly reports an invalid error.

The following table describes closed enhancement requests:

Bug	Description
PLAT-19630	When you define a domain, you can enable password complexity to validate the password strength.
PLAT-18366	Client side errors are now handled by the Node process without initiating its shutdown.

Enterprise Data Catalog Fixed Limitations (10.2.1)

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

The following table describes fixed limitations:

Bug	Description
LDM-6588	An error appears when you run an SAP R/3 resource with the First N rows sampling option.
LDM-6555	When you run the profiling scanner on a Microsoft SQL Server data object, the Invalid utf8 data error appears.

Bug	Description
LDM-6215	When you run a Teradata resource after you choose a sampling option, the column profile results take a long time to appear. This issue occurs because Enterprise Data Catalog runs the column profile on full tables instead of a data source sample.
LDM-6109	When you run a resource with the Random sampling option, the error message SEVERE message "FAILED to fetch statistic appears in the Data Integration Service logs even when the profile run is successful.
LDM-5997	The Metadata Load Settings > Basic Profile Settings > Source Connection Name option does not display the correct tool tip.
LDM-5544	When you run a column profile on an XML data object, the data validation error appears if a column field name in the data object contains the - character.
EIC-3269	If you enable the Similarity Profile option and choose First N rows sampling option for a Microsoft SQL Server 2014 resource or Microsoft SQL Server 2016 resource that has more than 250 columns, and you run the resource on the Blaze engine, the similarity scanner fails with null pointer exception.
EIC-3152	Inconsistent profile results or value frequency results appear when you run a profile on Random N rows of a Microsoft SQL Server 2014 resource or Microsoft SQL Server 2016 resource on the Blaze engine.
EIC-2669	Similarity profile runs on all the rows in the data object even when you configure the sampling option as First N rows.

Enterprise Data Lake Fixed Limitations (10.2.1)

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

The following table describes fixed limitations:

Bug	Description
IDL-7995	When you apply a filter containing multibyte characters in the Data Preview tab, the preview does not return any data.
IDL-6679	When you upload a delimited file, if you set the Hive Table Storage Format property for the Enterprise Data Lake Service to Parquet or ORC and select a non-default schema to publish the asset to, Enterprise Data Lake creates a table for the asset in the incorrect schema.
IDL-6556	When you publish a worksheet in which a numeric column included in an aggregation calculation is inferred as a string, Enterprise Data Lake generates an invalid maplet.
IDL-6410	When you perform an aggregate calculation on a worksheet, the operation excludes columns that contain a single numeric value.
IDL-5531	The data preparation page for a project fails to open if the web browser you use to access Enterprise Data Lake does not have the self-signed certificates or certificate authority (CA)-signed certificates installed on the node hosting the Enterprise Data Lake Service.
IDL-5102	Data preview of a postgres table fails if the schema name consists of lower case characters.

Bug	Description
IDL-5101	Data preview of a MySQL table that includes the BIGINT UNSIGNED data type fails with the following error: Failed to preview data. Possible reasons include insufficient privileges to view the data. Contact your administrator or choose a different connection.
IDL-4987	If you log out of Intelligent Data Lake and log in, the Visualization tab does not appear in the data asset views.
IDL-4973	If the system directory is encrypted, upload to the Hive database fails.
IDL-4955	In a fully profiled data asset, decimal data appears as null values after the data asset is published.
IDL-4832	If the data asset contains ultibytes and extended ASCII characters, data preview and tabular preview show incorrect data.
IDL-1716	Publication of a worksheet that includes an inner join fails if Enterprise Data Lake infers string columns in the worksheet as Integers.
BDM-10566	If you create encrypted zones in the cluster and run the mappings or preview data in native mode, the mapping fails.

The following table describes closed enhancement requests:

Bug	Description
IDL-5514	Enterprise Data Lake supports LDAP authentication for HiveServer2.

Exception Management Fixed Limitations (10.2.1)

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

The following table describes fixed limitations:

Bug	Description
IDQ-5701	The Analyst tool displays an error when you update a cell to a null value in a cluster step. The issue arises when the column that contains the cell uses the varchar(1) data type.

Mappings and Workflows Fixed Limitations

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

The following table describes fixed limitations:

Bug	Description
PLAT-20528	A mapping fails when the mapping is configured with parallelism set to Auto and the Data Integration Service is configured with maximum parallelism set to a value greater than 1.
OCON-12314	Concurrent dynamic mappings display an incorrect error message in the mapping log when concurrency issues occur.
OCON-12105	Dynamic mappings fail intermittently with the following error message: [com.informatica.sdk.dtm.InvalidMappingException] Exception Message: [[LDTMEXP_0029] Failed to process mapping because of following reason [null]].].
MWF-1571	A workflow terminates when it tries to parse a system variable date.
MWF-1495	You can configure multiple sequence flows with the same name between two gateways in a workflow.
MWF-1493	If you run multiple concurrent workflows with Mapping tasks continuously over 24 hours in a hadoop environment, the Data Integration Service spawns an excessive number of threads and eventually stops.
MWF-1479	If you restart the workflow database, the Workflow Orchestration Service module on the Data Integration Service becomes unavailable.
MWF-1465	A workflow might fail to start when a mapping in the workflow contains a database connection that is not valid, even if the workflow does not call the database connection at run time.
MWF-1448	If a domain fails over from one gateway node to another gateway node while a workflow runs, the Data Integration Service re-runs any Mapping task that began before the domain failed. Any mapping that runs again fails during workflow recovery. The issue arises when the following conditions are true: <ul style="list-style-type: none"> - The workflow contains multiple parallel Mapping tasks and Command tasks between two Inclusive gateways. - The batch execution pool size specifies that the mappings in the Mapping tasks must run in two or more batches. - You manually recover the workflow or the workflow recovers automatically after the domain fails over. - You configure the nodes in a grid, and the grid Data Integration Service runs the workflow.
IDQ-5978	When an Exclusive gateway receives more than one default sequence flow from upstream gateways in a workflow, all of the default sequence flows that enter the gateway can run.
BDM-17539	When the Data Integration Service is disabled, the error message in the Developer tool incorrectly states that operating system profiles are disabled.
BDM-16777	If a master Data Integration Service loses contact with a worker Data Integration Service during a workflow run, the master service does not try to reconnect to the worker service. The master Data Integration Service defines the workflow status as failed.

Bug	Description
BDM-13914	A mapping that includes an Address Validator transformation can fail when the following conditions are true: <ul style="list-style-type: none"> - You set the Max Parallelism property on the mapping to Auto. - You set the Max Parallelism property on the Data Integration Service to a higher value than the Execution instances property on the transformation.
BDM-13067	A workflow fails during a Mapping task because workflow validation does not detect an unresolved parameter or connection in a mapplet that you embedded in a Match transformation.

Profiles and Scorecards Fixed Limitations (10.2.1)

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

The following table describes fixed limitations:

Bug	Description
IDE-3157	In the Analyst tool, drilldown does not work as expected for column profile results if you configured the 'ExecutionContextOptions.drilldown_page_size' custom property for Data Integration Service as n and if the number of rows returned by drilldown exceeds n.
IDE-3154	When the connection objects in the Informatica domain exceeds 500, column profiles take a long time to open in Informatica Developer.
IDE-3142	A column profile with a sampling option fails for Microsoft SQL Server data objects if the schema name contains the \$ character.
IDE-3095	When an operating system profile user runs a column profile on a flat file in Informatica Analyst, the The source data object was deleted. The profile results are from the last profile run. warning message appears if you configure the Source file directory as SourceDir (Parameter) for the flat file.
IDE-3056	Column profile run fails for a logical data object when the profile name exceeds 128 characters and you run the profile on the Hive engine in Hadoop run-time environment.
IDE-3052	When an operating system profile user runs a column profile, the profile run fails if the Use the Logged in User as Hadoop Impersonation User property for operating system profile user is set to true.
IDE-3028	In Informatica Developer, an exception appears when you export the results summary for an enterprise discovery profile.
IDE-3025	You can run profiles even when you do not have the Run Profiles and Scorecards privilege.
IDE-3009	Sometimes, the profile run fails when you run the enterprise discovery profile for Microsoft SQL Server schema objects in the Analyst tool.
IDE-2936	When you add a rule the second time to the scorecard, the previous scorecard run results display 0 rows.
IDE-2859	Profiling results do not appear for some column profiles in the Analyst tool but appears in the Developer tool.

Security Fixed Limitations (10.2.1)

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

The following table describes fixed limitations:

Bug	Description
PLAT-19421	When you modify the LDAP synchronization schedule in the Administrator tool, the Administrator might continue to display and use the former schedule.
BDM-9945	The Data Integration Service cannot use user impersonation to read data from an HDFS file that is located in an encryption zone.
BDM-14496	The Hive engine fails to run a Hive query due to a Kerberos authentication error if the query takes a long time to run.
BDM-12237	Kerberos authentication fails when you run a mapping that reads from an HDFS source and writes to an HDFS target on the Spark engine.
BDM-11116	When a mapping runs on the Blaze engine, the mapping fails while authenticating on Hive if the following conditions are true: <ul style="list-style-type: none">- Hiveserver2 is enabled with LDAP authentication.- The mapping reads from a Hive source and writes to a Hive target. The following error occurs: <code>Peer indicated failure: Error validating the log</code>
BDM-11092	Mappings fail when the mappings run longer than the Kerberos ticket expiration time on a Kerberos-enabled Hadoop cluster.

The following table describes closed enhancement requests:

Bug	Description
PLAT-18537	You can disable LDAP paging for groups in an LDAP security domain.
PLAT-16989	The domain log contains details on user lock out operations.

Third-Party Fixed Limitations (10.2.1)

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

The following table describes fixed limitations:

Bug	Description
BDM-17527	Mappings that run on Hadoop clusters that use Hive 1.x do not pad zeroes to match the scale of the column precision.. Reference number: HIVE-12063

Transformation Fixed Limitations (10.2.1)

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

The following table describes fixed limitations:

Bug	Description
OCON-13365	The REST Web Service Consumer transformation does not parse an unnamed array with a single instance of element.
IDQ-5965	If you run multiple identity match mappings concurrently through web services, the Data Transformation Manager (DTM) might fail to open the identity population data files.
IDQ-5651	If the Address Validator transformation reads a United States address that includes a character space between the digits in the house number, the mapping that contains the transformation stops unexpectedly.

Transformation Language Functions Fixed Limitations (10.2.1)

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

The following table describes fixed limitations:

Bug	Description
BDM-14555	The return value of an expression might be incorrect when the following conditions are true: <ul style="list-style-type: none">- The expression uses an argument that is a decimal port with scale 0.- The input to the decimal port is a constant.- The expression contains a function that honors precision and scale when evaluating arguments such as TO_CHAR or CAST.
BDM-10790	NULL data is loaded to the Hive target when you use a decimal with a precision less than 29 as an argument in the function RTRIM or LTRIM.

10.2.1 Known Limitations

Administrator Tool Known Limitations (10.2.1)

The following table describes known limitations:

Bug	Description
PLAT-20807	When you try to access search from the Administrator tool or open the error codes from the logs in Informatica Knowledge Base, the following link is not accessible: https://csm.informatica.com/informaticaCSM/infa9/csm/ACProxy/oktaLogin.jsp Workaround: Use the following Network link to access search from the Administrator tool or open the error codes from the logs in Informatica Knowledge Base: https://network.informatica.com/login!input.jspx?referer=https://csm.informatica.com
OCON-1138	When you import or export data through Sqoop, the Administrator tool does not display the correct execution statistics in the Monitoring tab. Workaround: See the execution statistics in the yarn log. (452798)

Analyst Tool Known Limitations (10.2.1)

The following table describes known limitations:

Bug	Description
ANT-682	When you try to delete an asset that another user changed, the Analyst tool fails to warn you that the asset is not the latest copy. (396636)
ANT-595	After you rename a domain, you cannot export a mapping specification to PowerCenter. Workaround: Use the Developer tool to export the mapping object.
ANT-45	When you try to export the mapping specification to PowerCenter Repository using pcclientsmartuser that run on a Windows network using two-factor authentication, the mapping specification export fails. Workaround: Export the mapping specification to PowerCenter Repository using INFAKRB.INFADEV.COM (460405)

Application Service Known Limitations (10.2.1)

The following table describes known 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.
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.

Bug	Description
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-1571	When you enable the Model Repository Service for DEBUG level of logging, the service generates multiple ERROR level log messages.
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-1462	You cannot restore or view the older versions of Model repository objects after you restore the contents of a Subversion repository to another Subversion repository.
BDM-4669	The Data Integration Service does not apply the cost-based optimization method when you configure the mapping to use load order constraints with the full optimizer level. (431534)
BDM-2483	The Processes tab of the Email Service includes an environment variable section even though environment variables are not supported for the Email Service. If you add an environment variable, the Email Service ignores it. (442102)
BDM-1828	If you run web service requests on a Data Integration Service grid and you incorrectly configure the external HTTP load balancer to use nodes with the service role only, the Data Integration Service does not redirect requests to nodes with both the service and compute roles. Some web service requests dispatched to the node with the service role only might fail. Workaround: Configure the external HTTP load balancer to use nodes with both the service and compute roles. (427052)
BDM-1798	When you run a mapping on a Data Integration Service grid configured to run jobs in separate remote processes, the Monitor tab of the Administrator tool might indefinitely list the mapping state as Running even though the infacmd command line program and the mapping log indicate that the mapping failed. (432316)
BDM-19090	The Data Integration Service hangs when the service runs concurrent SQL data service mappings.
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.

Bug	Description
BDM-13006	<p>Mappings run on a Data Integration Service grid might hang indefinitely when the following conditions are true:</p> <ul style="list-style-type: none"> - The Data Integration Service is configured to run jobs in separate remote processes. - The Resource Manager Service becomes unavailable after the Data Integration Service has been enabled and has elected a master compute node. <p>Workaround: Enable the Resource Manager Service to continue running the mappings. (439628)</p>
BDM-13004	<p>The DTM process does not create DTM log files for mappings included in workflow Mapping tasks when the following conditions are true:</p> <ul style="list-style-type: none"> - The Data Integration Service is configured to run jobs in separate remote processes. - The mapping included in the workflow Mapping task uses multibyte characters. <p>(443052)</p>
BDM-13002	<p>A Data Integration Service grid configured to run jobs in separate remote processes does not use a secure connection to communicate with remote DTM processes even though the domain is enabled for secure communication. (432752)</p>
BDM-12991	<p>In a Kerberos domain, mappings fail to run on a Data Integration Service grid configured to run jobs in separate remote processes.</p> <p>Workaround: Configure the Data Integration Service to run jobs in separate local processes. (435471)</p>
BDM-12990	<p>When you update the compute role on a node assigned to a Data Integration Service grid and then recycle the Data Integration Service, you might encounter inconsistent behavior across the Informatica client tools. For example, mappings might fail to run in the infacmd command line program but succeed in the Developer tool.</p> <p>Workaround: Restart the domain. (436753)</p>
BDM-12987	<p>The consolidated log file for a mapping might contain the incorrect DTM log file when the following conditions are true:</p> <ul style="list-style-type: none"> - The Data Integration Service is configured to run jobs in separate remote processes. - The Mapping task in a workflow is configured to save the Mapping task log file by the number of Mapping task runs. <p>Workaround: Configure the Mapping task to save the Mapping task log file by timestamp. (439632)</p>

Big Data Management Known Limitations (10.2.1)

The following table describes known limitations:

Bug	Description
OCON-9377	<p>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.</p>
OCON-9376	<p>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.</p>
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>

Bug	Description
OCON-9005	When you run TDCH mappings on the Hive engine to write time data to a Teradata target, the nanosecond part is truncated.
OCON-8850	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.
OCON-8779	If you configure Sqoop to export data of the Real data type to IBM DB2 z/OS targets, the mapping fails.
OCON-7687	When you export data through Sqoop and the columns contain mixed case characters, the mapping fails.
OCON-7669	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. Workaround: Use the generic Oracle JDBC driver to export data.
OCON-7429	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. Workaround: Use the <code>hdp-connector-for-teradata-1.4.1.2.3.2.0-2950-distro.tar.gz</code> JAR.
OCON-730	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)
OCON-7291	Mappings that read data from a Teradata source and contain the != (not equal) operator in the filter override query fail. This issue occurs if you run the Teradata Parallel Transporter mapping on a Hortonworks cluster and on the Blaze engine. Workaround: Use a native expression with the ne operator instead of the != operator.
OCON-7280	If you configure Sqoop and update the columns in the advanced SQL query, the mapping fails on the Blaze engine.
OCON-7216	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.
OCON-7212	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. Workaround: Before you run the mapping, create a table in the target database with columns corresponding to the connected ports.
OCON-7205	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.
OCON-7078	Sqoop mappings that import data from or export data to an SSL-enabled database fail on the Blaze engine.
OCON-7076	When you run a Sqoop mapping and abort the mapping from the Developer tool, the Sqoop map-reduce jobs continue to run. Workaround: On the Sqoop data node, run the following command to kill the Sqoop map-reduce jobs: <code>yarn application -kill <application_ID></code>

Bug	Description
OCON-688	When you enable Sqoop for a logical data object and export data to an IBM DB2 database, the Sqoop export command fails. However, the mapping runs successfully without any error. (456455)
IN-3362	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.
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: SEVERE: UserDefinedJob execution failed
BDM-9987	A mapping configured with a filter expression upstream of a Joiner transformation might be inconsistent with map-side joins.
BDM-6389	A mapping fails to add statistics to Hive table metadata after loading data to the table on Hortonworks. Workaround: To view statistics for a table, run the following command on the HIVE command line: ANALYZE TABLE <table name> COMPUTE STATISTICS;
BDM-4597	A mapping with a joiner transformation that processes more than 4,294,967,294 rows in a single partition will fail. Workaround: If possible, increase partitioning on the source.
BDM-4107	If a mapping or workflow contains a parameter, the mapping does not return system-defined mapping outputs when run in the Hadoop environment.
BDM-3853	When the Blaze engine runs a mapping that uses source or target files in the WASB location on a cluster, the mapping fails with an error like: java.lang.RuntimeException: [<error_code>] The Integration Service failed to run Hive query [exec0_query_6] for task [exec0] due to following error: <error_code> message [FAILED: ... Cannot run program "/usr/lib/python2.7/dist-packages/hdinsight_common/decrypt.sh": error=2, No such file or directory], ... The mapping fails because the cluster attempts to decrypt the data but cannot find a file needed to perform the decryption operation. Workaround: Find the following files on the cluster and copy them to the /usr/lib/python2.7/dist-packages/hdinsight_common directory on the machine that runs the Data Integration Service: - key_decryption_cert.prv - decrypt.sh
BDM-3687	When you run a Sqoop mapping on the Spark engine, the Sqoop map-reduce jobs run in the default yarn queue instead of the yarn queue that you configure. Workaround: To run a map-reduce job in a particular yarn queue, configure the following property in the Sqoop Arguments field of the JDBC connection: -Dmapreduce.job.queueName=<NameOfTheQueue> To run a Spark job in a particular yarn queue, configure the following property in the Hadoop connection: spark.yarn.queue=<NameOfTheQueue>
BDM-2222	The Spark engine does not run the footer row command configured for a flat file target. (459942)

Bug	Description
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-17849	<p>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:</p> <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</pre> <p>[ICMD_10033] Command [ExportObjects] failed with error [[OIECMD_10006] Error occurred during export function. See previous error.].</p> <p>Workaround:</p> <ol style="list-style-type: none"> Copy the relaxngDatatype-20020414.jar file and the xsom-20100725.jar file from the following locations: <ul style="list-style-type: none"> Developer Client: DeveloperClient\connectors\thirdparty\infa.xmlparser\common Server: \$INFA_HOME/connectors/thirdparty/infa.xmlparser/common/ Place the JAR files in the following locations: <ul style="list-style-type: none"> Developer Client: DeveloperClient\infacmd\plugins\tools\eclipse\plugins Server: \$INFA_HOME/isp/bin/plugins/tools/eclipse/plugins Go to the configuration folder and delete all the files except the config.ini file. <ul style="list-style-type: none"> Developer Client: DeveloperClient\infacmd\plugins\tools\eclipse\configuration Server: \$INFA_HOME/isp/bin/plugins/tools/eclipse/configuration
BDM-17755	<p>When a decimal port on a partitioned, bucketed Hive target is disconnected, a passthrough mapping that runs on the Blaze engine fails.</p> <p>Workaround: Reconnect all the ports in the target.</p>
BDM-17685	<p>The Data Integration Service does not remove temporary empty files from the operating system profile target directory when all the following conditions are true:</p> <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. <p>Workaround: Manually delete the temporary empty files from the target directory that you configured for the operating system profile.</p>
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-17481	If the Data Integration Service is enabled to use operating system profiles, mappings with a Lookup transformation fail in Hive mode on a MapR cluster.

Bug	Description
BDM-17362	<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.
BDM-17330	<p>The Spark engine creates Hive tables in Azure HDInsight and Hortonworks clusters even if the impersonation user does not have privileges to create a table.</p>
BDM-17264	<p>A mapping that reads a large number of reference tables may take longer than expected to run on the Hive engine. The issue is observed when the mapping includes transformations that collectively read 140 reference tables.</p> <p>Workaround: Run the mapping on the Blaze engine.</p>
BDM-17175	<p>After you import the workflow into the Model repository, you cannot create a parameter set for the SQL parameter with the empty default value.</p>
BDM-17174	<p>When memory usage reaches the maximum container size, YARN kills the container.</p> <p>Memory usage on the OOP Container Manager reaches the maximum container size if the following conditions are true:</p> <ul style="list-style-type: none"> - Concurrent jobs run for longer than two days. - The Blaze engine has not reached the idle timeout limit or the sunset time.
BDM-17144	<p>When you edit the reusable Lookup transformation properties in the mapping level before importing it into the Model repository, the Developer tool imports a non-reusable Lookup transformation without assigning a connection.</p> <p>Workaround: Assign a connection manually to the Lookup relational data object or instance after import.</p>
BDM-17141	<p>When you have a PowerCenter mapping with multiple sources pointing to the same Source Qualifier, and import it into the Developer tool, no connections are assigned to the objects with multiple sources.</p> <p>Workaround: Manually assign a connection in the mapping level.</p>
BDM-17110	<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>
BDM-17101	<p>When you run high concurrency mappings on the Spark engine on an Amazon EMR cluster, some mappings fail with an error like "FileNotFoundException: File does not exist: hdfs://<path>/spark_conf.zip".</p> <p>Workaround: You can use cluster workflows to divide the processing workload among more than one cluster.</p>
BDM-16906	<p>You cannot provide different parameter values for multiple flat file source instance properties in the session level that points to the same data object parameter.</p> <p>Workaround: Clone the physical data object and assign the required parameter values for the source instance.</p>
BDM-16843	<p>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.</p>

Bug	Description
BDM-16521	<p>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.</p> <p>Workaround: Edit the column type in the Lookup transformation to run the mapping.</p>
BDM-14626	<p>When you do not assign a connection in the physical data object, the workflow application fails to deploy.</p> <p>Workaround: After importing the mapping, assign a connection to the physical data object and mapping. Or, assign a connection in the import wizard in the Developer tool.</p>
BDM-14139	<p>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:</p> <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. <p>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.</p>
BDM-13107	<p>When you run a mapping with Spark monitoring enabled, performance bottlenecks can appear. The performance varies based on the mapping complexity.</p>
BDM-12997	<p>Mappings that read from one of the following sources fail to run in the native environment when the Data Integration Service is configured to run jobs in separate remote processes:</p> <ul style="list-style-type: none"> - Flat file or complex file in the Hadoop Distributed File System (HDFS) - HIVE table - HBase table <p>Workaround: On the Compute view for the Data Integration Service, configure the <code>INFA_HADOOP_DIST_DIR</code> environment variable for each node with the compute role. Set the environment variable to the same value configured for the Data Integration Service Hadoop Distribution Directory execution option for the Data Integration Service. (443164)</p>
BDM-1271	<p>If you define an SQL override in the Hive source and choose to update the output ports based on the custom query, the mapping fails on the Blaze engine.</p>
BDM-10895	<p>If the JSON object that you are reading from has a key name with a dot operator (.) and you use the key name in an expression, the following error occurs: <code>Cannot read complex definition</code></p>
BDM-10570	<p>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.</p> <p>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.</p>

Bug	Description
B2BT-362	<p>A complex file data object containing an intelligent structure model that parses Microsoft Excel or Microsoft Word might fail in a Cloudera CDH 5.11 or CDH 5.12 environment.</p> <p>Workaround: For a CDH 5.11 and CDH 5.12 environment, copy the following files to the directory <code>\$INFA_HOME/services/shared/hadoop/CDH_5.10/spark/jars/</code>:</p> <ul style="list-style-type: none"> - <code>\$INFA_HOME/services/shared/jars/thirdparty/org.apache.poi-poi-ooxml-3.17.jar</code> - <code>\$INFA_HOME/services/shared/jars/thirdparty/org.apache.poi-poi-3.17.jar</code> - <code>\$INFA_HOME/services/shared/jars/thirdparty/org.apache.commons-commons-collections4-4.1.jar</code> - <code>\$INFA_HOME/services/shared/jars/thirdparty/org.apache.poi-poi-ooxml-schemas-3.17.jar</code> - <code>\$INFA_HOME/services/shared/jars/thirdparty/org.apache.xmlbeans-xmlbeans-2.6.0.jar</code> <p>After you copy the files, restart the Data Integration Service.</p>
B2BT-333	<p>A complex file data object containing an intelligent structure model that parses Microsoft Excel or Microsoft Word might fail in a MapR environment.</p> <p>Workaround: For a MapR environment, copy the following files to the directory <code>\$INFA_HOME/services/shared/hadoop/MapR_6.0/spark/jars/</code>:</p> <ul style="list-style-type: none"> - <code>\$INFA_HOME/services/shared/jars/thirdparty/org.apache.poi-poi-ooxml-3.17.jar</code> - <code>\$INFA_HOME/services/shared/jars/thirdparty/org.apache.poi-poi-3.17.jar</code> - <code>\$INFA_HOME/services/shared/jars/thirdparty/org.apache.commons-commons-collections4-4.1.jar</code> - <code>\$INFA_HOME/services/shared/jars/thirdparty/org.apache.poi-poi-ooxml-schemas-3.17.jar</code> - <code>\$INFA_HOME/services/shared/jars/thirdparty/org.apache.xmlbeans-xmlbeans-2.6.0.jar</code> <p>After you copy the files, restart the Data Integration Service.</p>

Big Data Streaming Known Limitations (10.2.1)

The following table describes known limitations:

Bug	Description
IIS-924	When you project columns in XML format, the Streaming mapping fails if the XSD schema contains definitions for namespace and references.
IIS-920	When you read XML payloads in Streaming mappings, the default values specified in the XSD schema are not read by the StaxXMLParser and StaxXMLGenerator.
IIS-919	When you read XML payloads in Streaming mappings, the nillable constraint specified in the XSD schema is not read by the StaxXMLParser and StaxXMLGenerator.
IIS-914	You cannot specify attributes for simple elements in XML payloads in Streaming mappings.
IIS-1771	The Administrator tool does not display monitoring statistics for applications and streaming mappings that run on ticketed clusters on a MapR distribution.

Bug	Description
IIS-1736	When you use a JDBC target in a streaming mapping to write mixed case column titles to an Oracle table, the data is not written to the table. Workaround: To write mixed case column titles to an Oracle table, create the table by running DDL statements from an SQL client.
IIS-1735	When you write to a JDBC target in a streaming mapping where the fields in the JDBC target are nullable by selecting the flag, the mapping fails. Workaround: To write nullable values to the JDBC target, create the table by running the DDL statements from an SQL client.
IIS-1476	In a streaming mapping, when you use a data masking transformation on source metadata of binary data type, the application fails with a fatal error.
IIS-1138	Streaming mappings that read from MapR Streams sometimes fail with the following error: Failed to get records

Domain Known Limitations (10.2.1)

The following table describes known limitations:

Bug	Description
PLAT-20922	When you try to edit a password, a message with the rules to set a complex password appears even if you have disabled password complexity. Workaround: You can ignore this message.

Developer Tool Known Limitations (10.2.1)

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: The transformation instance cannot be connected to the active transformations. 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.
OCON-9151	When you run a data preview on an Oracle table and the Number data type has a negative scale, data preview fails.

Bug	Description
OCON-6656	Mappings that contain columns of the Timestamp with Time Zone data type fail if you do not install the Microsoft Visual C++ 2010 Redistributable Package on the server machine.
OCON-609	When you import Teradata and Netezza mappings that contain stored procedures or SQL transformations from PowerCenter to the Developer tool, the import fails. (461184)
OCON-436	<p>When you import Teradata and Netezza mappings that contain parameters from PowerCenter into the Developer tool, the mapping conversion framework does not bind the parameters between the object and the mapping automatically.</p> <p>Workaround: When you import Teradata and Netezza mappings, mapping parameters <code><param></code> from PowerCenter are renamed as <code><param_mappingname></code> in the object level in the Developer tool.</p> <p>To bind the parameters with the mapping parameters, perform one of the following tasks:</p> <ul style="list-style-type: none"> - On the Parameter tab of the Netezza or Teradata source and target object, select the required parameters, and then select Expose as Mapping Parameter. - Select the required parameter in the Netezza or Teradata source or target object, click the Instance value, select Specify by Parameter from the list, browse, and then bind the object parameter with the required mapping parameters. <p>(455937)</p>
OCON-328	If a Teradata mapping contains an SQL query that reads decimal columns with precision less than 15, the mapping fails to run after you import the mapping from PowerCenter to the Developer tool. (459488)
OCON-13397	<p>After you upgrade and connect to the Model Repository Service from the Developer Client, relational connections in mappings from earlier versions might get unassigned.</p> <p>Workaround: Reconnect to the Model Repository Service or reopen the Developer Client.</p>
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.
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.
BDM-15525	When you run data preview on an Oracle table with a native SSL connection or you run a mapping that has an Oracle data object with a native SSL connection, the Developer tool shuts down unexpectedly. (393023)
BDM-14975	<p>When you use an SQL mapping and specify complex queries to generate a mapping, the Developer tool stops responding.</p> <p>Workaround: You must increase the default value of -Xmx value to 1536M in the <code>developerCore.ini</code> file and relaunch the Developer tool. (458862)</p>

Enterprise Data Catalog Known Limitations (10.2.1)

The following table describes known limitations:

Bug	Description
EIC-4202	You cannot view the lineage information for flat file sources and targets of a PowerCenter resource on a Microsoft Windows machine and a File System resource. This issue occurred for a File System resource configured using the SMB/CIFS file protocol.
LDM-5217	You cannot terminate the scanner job if the scanner is in PauseRequested or CancelRequested state. Workaround: Kill the yarn job corresponding to scanner. If the scanner job does not display Failed state, then restart the Catalog Service.
LDM-5226	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.
EIC-4545	You cannot delete multi-valued custom attribute values of reference data type in Enterprise Data Catalog.
EIC-4632	You cannot delete custom attributes in Enterprise Data Catalog if you create custom attributes using the REST API.
LDM-5188	When you search for an asset name that contains an underscore (_), the search results page does not return any result.
LDM-6619	When you run the profiling scanner on the File System, HDFS, Azure Data Lake Store, OneDrive, SharePoint, Microsoft Azure Blob Storage, or Amazon S3 resources, an exception appears if the file name of the data source contains quotes.
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-4329	When you run a Netezza resource, the core scanner extracts junk tables as tables and the profiling scanner fails when it runs on the junk tables.
LDM-5081	When you delete an Informatica Axon resource, Enterprise Data Catalog does not delete the custom attribute assigned to the asset from the catalog.
LDM-5222	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-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.
EIC-3853	When Enterprise Data Catalog performs profile validation, the profiling scanner stops to run the profile on the remaining tables if the validation fails for the first table or second table.
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-4623	The profile validation for an Oracle resource fails with the following error: Profiling exception occurred: java.lang.RuntimeException: java.lang.RuntimeException: The input Type object is null

Bug	Description
EIC-4609	The detailed lineage in the lineage and impact view does not display the Java code in a java transformation for a PowerCenter resource.
EIC-1486	The Transformation Logic view displays inaccurate mapping names for the Informatica Cloud Services transformations.
LDM-3557	The File System resource fails with OutOfMemory error when scanning a large single JSON file.
LDM-3735	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.
LDM-4345	The <code>relationships</code> REST API returns different <code>levelCount</code> details for the same query.
EIC-4378	Sometimes, when you run multiple Informatica Platform resources at once, one of the Informatica Platform resource fails with the following error: Internal error. Secure communication failed because of the following SSL error: [SSLUtil_0004] The <code>infa_truststore.jks</code> file cannot be found. The truststore file must be in JKS format and must contain a certificate in PKCS12 format Workaround: Set the following JVM parameters in the Advanced properties section for the resource: -DINFA_TRUSTSTORE= <java trust store path> -DINFA_TRUSTSTORE_PASSWORD=<trust store encrypted key>
EIC-4306	Sometimes, Enterprise Data Catalog displays incorrect lineage information for an Informatica Platform resource if a workflow contains a mapping. Workaround: Run the resource again.
EIC-1837	Relationship diagram does not show synonyms.
EIC-4654	Profiling scanner fails for a JDBC resource if the data source contains an unsupported data type.
EIC-3848	Profiling scanner does not run on Avro and Parquet files if the files are located in OneDrive, SFTP, or remote Windows location.
LDM-6403	Incorrect profile results appear when you perform data domain discovery on compressed files.
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-2970	If a mapping contains an SQL or lookup transformation , Enterprise Data Catalog does not display the source qualifier and displays a source table that is not linked in the detailed lineage view for a PowerCenter resource Workaround: drill down on all the columns in the target.
LDM-5192	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.
EIC-3003	Erwin resource does not extract metadata from tables that share the same name.
LDM-5783	Enterprise Data Catalog runs the profiling scanner only on the first 1 MB of data for a JSON data source. The scanner ignores the rest of the file.

Bug	Description
EIC-4559	Enterprise Data Catalog does not display the lineage of a report when the connection links are not generated for the report.
EIC-3357	Enterprise Data Catalog does not display lineage information if you use a Database Script resource and the script contains an INSERT ALL query.
EIC-4640	Enterprise Data Catalog does not add a business title to a resource if the resource name contains a locale character.
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-2975	Duplicate mapping objects appear in the lineage and impact view for a table linked to a mapping containing the lookup transformation.
LDM-3912	Custom metadata resource does not extract connection parameters for Business Intelligence metadata sources.
EIC-4626	An Informatica Platform resource fails if a mapping contains unstructured files, such as PDF.
LDM-5220	After you fix the incorrect details in a reusable configuration used for a resource, the changes do not take effect on the resource.
EIC-4621	After you enable detailed lineage to configure and run an Informatica Platform resource with a target version earlier than version 10.2.1, the scanner runs successfully without an error but metadata is not extracted.
EIC-3547	After you create a mapping with a lookup transformation and connect a column to a table, the lineage and impact view displays incorrect summary and detailed lineage information for an Informatica Platform resource in Enterprise Data Catalog.
EIC-4651	After you create a PowerCenter mapping with a stored procedure transformation the connection links are missing between the stored procedure transformation and stored procedure in the database.
LDM-6668	<p>A run-time error appears for a data domain or composite data domain if the following conditions are true:</p> <ol style="list-style-type: none"> 1. You open the Catalog Administrator in Internet Explorer version 11. 2. You create and save a data domain or composite data domain. <p>This issue occurs only for data domains and composite data domains that have a non-EN character in their names.</p>
EIC-4776	<p>A run-time error appears in the Overview tab of an asset if the following conditions are true:</p> <ol style="list-style-type: none"> 1. Add an image in the Description section. 2. Save the description. 3. Edit and save the Description section.

Enterprise Data Lake Known Limitations (10.2.1)

The following table describes known limitations:

Bug	Description
IDL-8248	Because Enterprise Data Lake does not collect Hive statistics, row count statistics might not be accurately displayed for an asset created in Enterprise Data Lake using Blaze.
IDL-7196	Data preparation of a JSON Lines file fails if the file contains an unnamed array as the top level element.
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.

Parameters Known Limitations (10.2.1)

The following table describes known 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. <p>Workaround: Run, preview, or show the execution plan in the original mapping.</p>

Profiles and Scorecards Known Limitations (10.2.1)

The following table describes known limitations:

Bug	Description
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.
IDE-3115	Column profile fails for a Hive partitioned source.
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.

Bug	Description
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.
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.

Reference Data Known Limitations (10.2.1)

The following table describes known limitations:

Bug	Description
IDQ-6656	<p>You cannot import an XML object to the Model repository when the following conditions are true:</p> <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 Specification Known Limitations (10.2.1)

The following table describes known limitations:

Bug	Description
SYM-1428	If a rule specification contains multiple rule sets and you click on a rule set, the rule set might move to become a child of another rule set. The issue arises when the rule specification is editable and you zoom in and zoom out of the rule specification.
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.
IDQ-4250	The Analyst tool returns incorrect results when you test a rule specification that contains a mapplet that you generated from another rule specification in the same session. The issue arises when the mapplet that you generated reads another mapplet in the Model repository. Workaround: Log out of the Analyst tool and log back in. Ignore any error message that the Analyst tool displays.
IDQ-4249	The Analyst tool returns incorrect results when you test a rule specification that contains a mapplet that you generated from another rule specification. The issue arises when the rule specification that generated the mapplet contains a rule set with the same name as a mapplet in the Model repository. Workaround: Log out of the Analyst tool and log back in. Ignore any error message that the Analyst tool displays. (439453)
IDQ-4246	The Analyst tool might display an error message when you open a rule specification that contains a mapplet that you generated from another rule specification. The issue arises if you generate another version of the mapplet after you added the mapplet to the rule specification in the same Analyst tool session. Workaround: Log out of the Analyst tool and log back in. Ignore any error message that the Analyst tool displays. (439258)

Security Known Limitations (10.2.1)

The following table describes known 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 Known Limitations (10.2.1)

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
OCON-7505	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. Cloudera ticket reference number: 124305
OCON-7504	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. Cloudera ticket reference number: 124302
OCON-7503	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. Cloudera ticket reference number: 124306
OCON-7459	When you export data to an IBM DB2 target through Sqoop, the mapping fails if all of the following conditions are true: <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. Apache ticket reference number: SQOOP-3212

Bug	Description
OCON-7431	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. Cloudera ticket reference number: 124302
OCON-7219	When you run a Sqoop mapping on the Blaze engine to export Teradata float data, the data is truncated after the decimal point. Cloudera support ticket number: 113716
OCON-7214	Sqoop mappings fail on the Blaze engine if you use a custom query with the Order By clause to import data. Sqoop JIRA issue number: SQOOP-3064
OCON-7213	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. Sqoop JIRA issue number: SQOOP-2837
OCON-7211	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. Sqoop JIRA issue number: SQOOP-2349
OCON-2847	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) Data Direct case number: 00343832
OCON-1100	When you export data to an IBM DB2 z/OS database through Sqoop and do not configure the batch argument, the mapping fails. Workaround: Configure the batch argument in the mapping and run the mapping again. (459671) Apache ticket reference number: SQOOP-2980
BDM-9585	Mappings fail on the Spark engine when you configure an SQL override to access a Hive view. Apache Spark ticket reference number: SPARK-21154.
BDM-17470	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. Workaround: In the hive-site.xml on the cluster, set the hive.vectorized.execution.enabled property to false. Apache Hive ticket reference number: HIVE-14076
BDM-17204	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. HADOOP-12604
BDM-17020	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.

Bug	Description
BDM-16599	Concurrent mappings running on a Cloudera cluster fail when the KMS authentication service stops unexpectedly and the cluster becomes unusable. The cluster log contains a "Failed to renew token" error, and the session log contains a "Connection refused" error. The mapping succeeds again after you restart the KMS service. HADOOP-13420
BDM-14438	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. SPARK-22474
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-13650	Mappings fail on the Spark engine when the Spark engine runs on a secure HA cluster and the hadoop.security.auth_to_local property in the core-site.xml file contains a modified value. The mappings fail due to the following error: Failed to renew token Workaround: Configure the following property in the yarn-site.xml file using any node in the Hadoop cluster: yarn.resourcemanager.address=<node name> MapR ticket reference number: MAPREDUCE-6484
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
BDM-10410	For IBM BigInsights and Hortonworks clusters, insert into a Hive target table on Amazon S3 can fail if the table is enabled for transactions.

Transformations Known Limitations (10.2.1)

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-6912	When you add a mapplet to a Match transformation, the Match Fields option becomes inactive when the following conditions are true: <ul style="list-style-type: none"> - You add one or more parameters to the transformations in the mapplet. - You create mapping parameters that specify values for properties on the Match transformation. - You resolve the mapping parameters in the mapping canvas.

Bug	Description
IDQ-6664	<p>A mapping that contains a Labeler transformation or a Parser transformation fails validation when the following conditions are true:</p> <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. <p>When the conditions are true, the probabilistic model no longer appears in the transformation strategy.</p>
IDQ-6662	<p>A mapping that contains a Classifier transformation fails validation when the following conditions are true:</p> <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. <p>When the conditions are true, the classifier model no longer appears in the transformation strategy.</p>

Workflows Known Limitations (10.2.1)

The following table describes known limitations:

Bug	Description
MWF-1430	<p>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.</p>
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>

Informatica Global Customer Support

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

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

<http://www.informatica.com/us/services-and-training/support-services/global-support-centers>.

If you are an Informatica Network member, you can use Online Support at

<http://network.informatica.com>.