



Contents

| | |
|--|----|
| Preface. | 2 |
| Support Changes. | 2 |
| Technical Preview Lifted. | 2 |
| Verify System Requirements. | 3 |
| Installation and Upgrade. | 3 |
| Upgrade Path. | 3 |
| Service Pack Files. | 3 |
| Download the Service Pack Files. | 4 |
| Update the Input.properties File. | 4 |
| Run the Installer. | 5 |
| Post-installation Steps. | 6 |
| Post-installation Steps for the Analyst Service. | 6 |
| Post-installation Tasks for Enterprise Data Catalog. | 6 |
| Post-Installation Steps for the Python Transformation. | 6 |
| Emergency Bug Fixes Merged into 10.4.1.1. | 8 |
| 10.4.1.1 Fixed Limitations and Closed Enhancements. | 8 |
| Data Engineering Integration Fixed Limitations (10.4.1.1). | 8 |
| Data Engineering Streaming Fixed Limitations (10.4.1.1). | 9 |
| Data Privacy Management Fixed Limitations (10.4.1.1). | 10 |
| Enterprise Data Catalog Fixed Limitations and Closed Enhancements (10.4.1.1). | 11 |
| Enterprise Data Preparation Fixed Limitations (10.4.1.1). | 13 |
| Security Fixed Limitations and Closed Enhancements (10.4.1.1). | 13 |
| Third-Party Fixed Limitations (10.4.1.1). | 14 |
| 10.4.1.1 Known Limitations. | 14 |
| Data Engineering Integration Known Limitations (10.4.1.1). | 15 |
| Data Engineering Streaming Known Limitations (10.4.1.1). | 15 |
| Data Privacy Management Known Limitations (10.4.1.1). | 16 |
| Enterprise Data Catalog Known Limitations (10.4.1.1). | 16 |
| Profiles and Scorecards Known Limitations (10.4.1.1). | 17 |

| | |
|--|----|
| Third-Party Known Limitations (10.4.1.1) | 18 |
|--|----|

This document contains important information about restricted functionality, known limitations, and bug fixes in Informatica version 10.4.1.1.

Preface

Informatica 10.4.1.1 is a service pack that contains multiple emergency bug fixes. The service pack supports Data Engineering, Data Security, and Data Catalog products. Traditional products, such as PowerCenter and Informatica Data Quality are not included in the service pack.

To get the service pack, download the 10.4.1.1 service pack from the [Informatica Network](#).

Support Changes

This section describes support changes in version 10.4.1.1.

Technical Preview Lifted

Effective in version 10.4.1.1, the following functionalities are lifted from technical preview:

Cloudera Data Platform Integration

For Data Engineering Integration, you can use Cloudera Data Platform (CDP) as a compute cluster to execute data engineering jobs in the Hadoop environment. You can use Cloudera CDP when you run data engineering jobs on the Spark engine. Cloudera CDP is not supported on the Blaze engine.

Cloudera CDP uses a base cluster and workload clusters to execute data engineering jobs. This architecture allows you to deploy workloads and share data among components by utilizing a shared catalog, unified security, consistent governance, and data life cycle management.

You can use Cloudera CDP when you run a mapping in the Hadoop environment with the following connections:

- PowerExchange for Amazon Redshift
- PowerExchange for Amazon S3
- PowerExchange for HDFS
- PowerExchange for Microsoft Azure Blob Storage
- PowerExchange for Microsoft Azure CosmosDB SQL API
- PowerExchange for Microsoft Azure Data Lake Storage Gen1
- PowerExchange for Microsoft Azure Data Lake Storage Gen2
- PowerExchange for Microsoft Azure SQL Data Warehouse
- Sqoop

You can also use Cloudera CDP when you run a PowerExchange for HDFS mapping in the native environment.

For more information, see the *Informatica® Data Engineering 10.4.1 Integration Guide*.

Verify System Requirements

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

In each release, Informatica can add, defer, and drop support for the non-native distributions and 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 Customer Portal:

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

Installation and Upgrade

Upgrade Path

Version 10.4.1.1 is a service pack that you apply on Informatica 10.4.1.

If you are on a version between 10.0.0 and 10.2.2, you must first upgrade to Informatica 10.4.1 and then apply Informatica 10.4.1.1.

If you are on version 10.4.0, 10.4.0 service pack, or a 10.4.0.2 cumulative patch, you must first apply Informatica 10.4.1, and then apply Informatica 10.4.1.1.

If you are on version 10.4.1 cumulative patch, apply Informatica 10.4.1.1 on Informatica 10.4.1 cumulative patch.

Note: The Informatica 10.4.0.2 and 10.4.1 cumulative patches are available for Enterprise Data Catalog.

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

Service Pack Files

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

The service pack includes the following files:

Input.properties

Identifies the root directory of the Informatica installation to which you will install the service pack. You update the file with the directory path. The file also contains a rollback property that you can set if you decide to uninstall the service pack.

install.bat

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

install.sh

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

Download the Service Pack Files

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

To apply the service pack on Informatica 10.4.1, you can download the service pack installer, Informatica Developer installer, and command line utilities package.

Download one of the following packages:

- `informatica_10411_server_linux-x64.tar`
Contains updates for Redhat Enterprise Linux installations.
- `informatica_10411_server_suse11-x64.tar`
Contains updates for SUSE Linux Enterprise Server installations.

Download the following package for Informatica Developer:

- `informatica_10411_client_winem-64t.zip`
Contains updates for the Developer tool installation.

Download one of the following packages for command line utilities:

- `informatica_10411_cmd_utilities_linux-x64.zip`
Contains updates for Redhat Enterprise Linux installations.
- `informatica_10411_cmd_utilities_suse11-x64.zip`
Contains updates for SUSE Linux Enterprise Server installations.

Update the Input.properties File

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

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

When you install or roll back the service pack, the installer applies all the components, by default. To apply or roll back a specific component, remove the comment tag (`#`) associated with the

component that you want to apply. For Data Engineering component, set BDM_ONLY to 1. For Enterprise Data Catalog, set EDC_ONLY to 1. For Enterprise Data Preparation, set EDP_ONLY to 1.

Data Privacy Management requires Data Engineering and Enterprise Data Catalog components. For Data Privacy Management, remove the comment tag(#) associated with Data Engineering, Enterprise Data Catalog, and Data Privacy Management, and set the values to 1.

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

1. Data Engineering products
2. Enterprise Data Catalog
3. Enterprise Data Preparation
4. Data Privacy Management

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

1. Enterprise Data Preparation
2. Enterprise Data Catalog

Note: If you had configured Data Asset Analytics, run the 10411_rollback.sql script on the database server after you roll back Enterprise Data Catalog. The script is available at the following location: <INFA_HOME>/services/CatalogService/DAABackupScripts/<database type>.

3. Data Engineering products
4. Data Privacy Management
5. Save and close the file.

Run the Installer

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

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

Note: After you apply version 10.4.1.1, the following message might appear in the `ServicePack_10.4.1.1_Install.log` file:

Cannot perform the <delete or backup> operation because the file does not exist: <file path>

You can ignore the message.

Post-installation Steps

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

Post-installation Steps for the Analyst Service

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

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

Post-installation Tasks for Enterprise Data Catalog

After you install Enterprise Data Catalog, complete the listed post-installation tasks.

Data Asset Analytics

After you apply 10.4.1.1 on 10.4.1.0.1, perform the following steps to synchronize events related to data asset enrichment and collaboration:

1. In Informatica Administrator, add the `LdmCustomOptions.ingest.store.events.on.reindex.bool` custom property for the Catalog Service and set the value as `true`.
2. Re-index the Catalog Service.

Download SAP Transports

Before you run the SAP BW, SAP BW/4HANA, and SAP S4/HANA scanners, download the compatible versions of the SAP transports and import them into the SAP server. The transports are located in the `SAP_Scanner_Binaries.zip` file that is downloaded from the Informatica installer location.

Post-Installation Steps for the Python Transformation

To use the Python transformation, you must ensure that the worker nodes on the Hadoop cluster contain an installation of Python after you install or upgrade.

Note: If you previously installed Python in the directory `<Informatica installation directory>/services/shared/spark/python`, you must reinstall Python.

Complete the different tasks depending on the product that you use.

Installing Python for Data Engineering Integration

To use the Python transformation in a mapping, the worker nodes on the cluster must contain a uniform installation of Python. You can ensure that the installation is uniform in one of the following ways:

Verify that the Python installation exists.

Verify that all worker nodes on the cluster contain an installation of Python in the same directory, such as `/usr/lib/python`, and that each Python installation contains all required modules.

Additionally, verify that the following Spark advanced property in the Hadoop connection is configured based on the directory that stores the Python installation:

```
infaspark.pythontx.executorEnv.PYTHONHOME
```

Install Python.

Install Python on every Data Integration Service machine. You can create a custom installation of Python that contains specific modules that you can reference in the Python code. When you run mappings, the Python installation is propagated to the worker nodes on the cluster.

If you choose to install Python on the Data Integration Service machines, complete the following tasks:

1. Install Python.
2. Optionally, install any third-party libraries such as numpy, scikit-learn, and cv2. You can access the third-party libraries in the Python transformation.
3. Copy the Python installation folder to the following location on the Data Integration Service machine:

```
<Informatica installation directory>/services/shared/spark/python
```

Note: If the Data Integration Service machine already contains an installation of Python, you can copy the existing Python installation to this location.

Changes take effect after you recycle the Data Integration Service.

Installing Python for Data Engineering Streaming

To use the Python transformation in a streaming mapping, you must install Python and the Jep package. Because you must install Jep, the Python version that you use must be compatible with Jep. You can use one of the following versions of Python:

2.7
3.3
3.4
3.5
3.6

To install Python and Jep, complete the following tasks:

1. Install Python with the **--enable-shared** option to ensure that shared libraries are accessible by Jep.
2. Install Jep. To install Jep, consider the following installation options:
 - Run `pip install jep`. Use this option if Python is installed with the pip package.

- Configure the Jep binaries. Ensure that `jep.jar` can be accessed by Java classloaders, the shared Jep library can be accessed by Java, and Jep Python files can be accessed by Python.
3. Optionally, install any third-party libraries such as numpy, scikit-learn, and cv2. You can access the third-party libraries in the Python transformation.
 4. Copy the Python installation folder to the following location on the Data Integration Service machine:


```
<Informatica installation directory>/services/shared/spark/python
```

Note: If the Data Integration Service machine already contains an installation of Python, you can copy the existing Python installation to this location.

Changes take effect after you recycle the Data Integration Service.

Emergency Bug Fixes Merged into 10.4.1.1

Informatica merged Emergency Bug Fixes (EBFs) from previous releases into version 10.4.1.1. These EBFs provided fixes for issues that were found in previous releases.

For a list of EBFs that were merged into version 10.4.1.1, see the following Informatica Knowledge Base article: <https://kb.informatica.com/faq/7/Pages/26/631029.aspx>

10.4.1.1 Fixed Limitations and Closed Enhancements

Data Engineering Integration Fixed Limitations (10.4.1.1)

Fixed Limitations

The following table describes fixed limitations:

| Issue | Description |
|-----------|---|
| BDM-34680 | When the Spark engine runs a mapping that executes a Hive query, the mapping might fail with error <code>java.lang.OutOfMemoryError: Java heap space</code> . |
| BDM-34734 | <p>On a Hortonworks HDP cluster, when the Blaze engine runs a mapping that reads date type data from a Hive source stored in Parquet format, the mapping fails.</p> <p>On an HDP 2.6 cluster, the mapping fails with the following error message: <code>[GRIDDTM_1016] The Integration Service failed to execute grid mapping with following error [An internal exception occurred with message: java.lang.ClassCastException: org.apache.hadoop.io.IntWritable cannot be cast to org.apache.hadoop.hive.serde2.io.DateWritable]</code></p> <p>On an HDP 3.1 cluster, the mapping fails with the following error message: <code>[GRIDDTM_1016] The Integration Service failed to execute grid mapping with following error [An internal exception occurred with message: java.lang.ClassCastException: org.apache.hadoop.io.IntWritable cannot be cast to org.apache.hadoop.hive.serde2.io.DateWritableV2]</code></p> |
| BDM-34445 | When the Blaze engine runs a mapping where the source is a Hive transactional table and column names in the Hive source have a case mismatch with columns in the physical data object of the mapping, the mapping might fail or return no data. |

| Issue | Description |
|------------------------|---|
| BDM-34343 BDM-31535 | Mappings fail intermittently in Spark mode when the Data Integration Service fails to fetch delegation token for Hive metastore. For more information, see the following Knowledge Base article: https://kb.informatica.com/solution/23/Pages/69/568300.aspx |
| BDM-33894 | When the Spark engine runs a mapping on a Cloudera CDP cluster with dynamic resource allocation enabled, the mapping fails with the following SPARK_1003 error: "Failed to connect to <node running Spark shuffle service>" |
| BDM-33751 BDM-34638 | When the Blaze engine reads a modified Hive partition table of text format, the source read fails with the following error: [CORE_3] An internal exception occurred with message: java.lang.RuntimeException: No data found for parttion state java.lang.RuntimeException: No data found for partition state |
| BDM-33677 | On a Hortonworks HDP 3.1 cluster, when the Blaze engine processes the TO_CHAR function to convert timestamp data from a Hive source to character strings, the mapping fails with the following error: SEVERE: The Integration Service failed to execute the mapping. java.lang.RuntimeException: java.util.concurrent.CompletionException: com.informatica.sdk.dtm.ExecutionException: [GRIDDTM_1016] The Integration Service failed to execute grid mapping with following error [An internal exception occurred with message: java.lang.ClassCastException: org.apache.hadoop.hive.serde2.io.TimestampWritableV2 cannot be cast to org.apache.hadoop.hive.serde2.io.TimestampWritable]. |
| BDM-31896 | When you run a mapping on a Hortonworks HDP 3.x cluster, the mapping fails when the following conditions are true: <ul style="list-style-type: none"> - In the Hadoop cluster properties, hive.server2.enable.doAs is set to false. - In the core-site.xml properties, hadoop.proxyuser.<proxy user>.groups and hadoop.proxyuser.<proxy user>.users do not include the Hive user. The mapping fails with the following error: org.apache.hadoop.ipc.RemoteException (org.apache.hadoop.security.authorize.AuthorizationException): User: <user> is not allowed to impersonate hive |

Data Engineering Streaming Fixed Limitations (10.4.1.1)

The following table describes fixed limitations:

| Bug | Description |
|----------|--|
| IIS-4827 | Data loss might occur for some of the targets when you run a streaming mapping with a single JMS source and multiple targets after you configure durable subscriptions. Only one of the multiple targets receives the data. |
| IIS-4815 | When dynamic schema strategy is enabled in a streaming mapping that has a Kafka source and Complex File data object target with no column projection, the target files are not moved to the target directory. This issue occurs when the executor goes to a dead state. The target files are moved as part of cleanup. |

| Bug | Description |
|----------|---|
| IIS-4813 | When you run a streaming mapping with a Microsoft Azure Data Lake Storage Gen2 or HDFS target that has Parquet payload on the Hortonworks HDP version 3.1.5 distribution, the target files are not moved to the target directory after clean up when the mapping is stopped and the log file does not record the error. |
| IIS-4811 | When you run a streaming mapping with an HDFS target that has Parquet payload on an unsecured Cloudera CDH cluster, the target files are not moved to the target directory after clean up when the mapping is stopped. |

Data Privacy Management Fixed Limitations (10.4.1.1)

The following table describes fixed limitations:

| Bug | Description |
|------------|--|
| SATS-30996 | You cannot use batch update in a data masking workflow in Test Data Management because the Data Privacy Management protection task does not include parameters for batch update. |
| SATS-30403 | If a user does not have the following two privileges, Data Privacy Management session shuts down unexpectedly: <ul style="list-style-type: none"> - View alerts - View data stores |
| SATS-30402 | Manual changes made to false negative information are not retained when you run a scan again. |
| SATS-30401 | A data store import job that includes location information fails. |
| SATS-30202 | A Reset Classification Policy job does not clear risk score details. |
| SATS-30199 | Editing a data domain changes the data domain group to SecureAtSourceService. |
| SATS-30193 | You require separate privileges to import catalog resources and to modify classification policies. |
| SATS-30191 | A scan that ends with a Warning status does not trigger a Reset Classification job. |
| SATS-30002 | A scan that includes a field that has more than 256 characters and qualifies for more than two data domains generates a row count error. |
| SATS-29938 | A subject registry scan on an Azure WASB data store goes into an infinite loop. |
| SATS-23135 | Resetting the classification results of a data store deletes the connection string and schema options of the data store. |
| SATS-23129 | If you integrate Axon Data Governance created with HTTPS with Data Privacy Management, sync jobs fail with a certificate error. |
| SATS-23127 | A security policy violation details page displays the wrong data stores and an incorrect data store count for data stores configured for Agent scans. |
| SATS-23110 | Repeat subject registry scans to link subjects on an unstructured data store return varied results. |

| Bug | Description |
|------------|---|
| SATS-21211 | A scan job that uses the remote agent fails if the scan includes a custom classification policy that has unsupported OOTB domains. |
| SATS-20790 | A scan that uses an agent calculates the confidence level based on all domains in the policy instead of only the domains included in the custom criteria. |
| SATS-15849 | User names appear with extra code values if you use LDAP authentication. |

Enterprise Data Catalog Fixed Limitations and Closed Enhancements (10.4.1.1)

The following table describes fixed limitations:

| Bug | Description |
|-----------|---|
| EIC-44464 | When you run a data domain discovery profile on a PostgreSQL source, the profiling service ignores the tables during the profile run if the source is configured as a JDBC resource. |
| EIC-43835 | The pre-script task fails for the multi-node cluster setup. |
| EIC-43661 | When you choose the native runtime environment and run a profile on a Microsoft Azure Data Lake Store Gen2 source, an exception appears if the resource has parquet files. |
| EIC-43582 | If you select a password protected package, the metadata scan of a File type SSIS resource fails. |
| EIC-43354 | The MicroStrategy resource does not create connection parameters at the database level. |
| EIC-43350 | You cannot delete or update a data domain in Catalog Administrator after you import the custom data domain from a .csv file. |
| EIC-43328 | When you run a profile on a PostgreSQL table, the profiling service does not run the profile on the columns with integer and bigserial data types. |
| EIC-43313 | The Catalog Service fails to start due to ingestion issues related to a File System resource. |
| EIC-43234 | The Informatica Platform resource fails to extract dynamic mapping metadata from an Informatica Platform data source when parsing a parameter file. |
| EIC-43163 | The Informatica Cluster Service fails to start with the following error when you use custom SSL certificates and remove the contents of the cluster periodically: <code>ERROR com.infa.products.ihs.beans.application.ClusterListener-com.infa.products.ihs.service.exception.InfaHadoopServiceException: [InfaHadoopServiceException_00084] Previous ssl-config-directory [/opt/informatica/keystore] exists. Please make sure infacmd ihs cleanCluster command has been used to clean up existing cluster.</code> |
| EIC-43112 | You cannot view the custom lineage for a Hive resource that is configured for case sensitivity. |
| EIC-43110 | The catalog fails to load the detailed lineage view for an SAP Business Objects resource. |

| Bug | Description |
|-----------|--|
| EIC-43085 | Unable to view the Data Asset Analytics tab with the following error message in the log files: Connection to the Data Asset Analytics repository failed. Make sure to verify the Data Asset Analytics configuration properties in Informatica Administrator or contact the Administrator for more information. The issue occurs if you use the DNS host name in the URL to access Enterprise Data Catalog. |
| EIC-4301 | You cannot view the asset type filter that is applied on the search results page after you apply additional filters from the Resources search tab. This issue occurs from the second page of the search results. |
| EIC-42897 | The File System resource fails to scan files in the sub directory on the remote Windows machine. |
| EIC-42896 | When you run a column profile on large string columns in Snowflake, Athena, or Exasol resources, the profile runs fails with an nsort error. |
| EIC-42754 | You cannot view lineage and impact information of views from a Snowflake resource. |
| EIC-42622 | When you choose Random N sampling option and run a column profile on an Oracle resource, the profile runs on all the rows if the resource uses a JDBC connection. |
| EIC-42621 | The detailed lineage view hangs unexpectedly for a SAP Business Objects resource. |
| EIC-42615 | Enterprise Data Catalog does not display the asset path for assets from the Netezza resource. |
| EIC-42589 | The catalog does not display the lineage and impact information of views after you run a Snowflake resource with tables and views. This issue occurs because the SQL DDL fails to parse. |
| EIC-42383 | When you choose the Run Discovery on both Source Metadata and Data option to run a data domain discovery profile, the catalog shows or accepts data domain recommendations if a column conforms to data rule or column rule. |
| EIC-42335 | If you do not specify the source directory during resource configuration, test connection for the Amazon S3 resource fails. |
| EIC-41832 | The Catalog Service fails to start after you apply version 10.4.0.2. |
| EIC-41645 | If you override the path of the target, the Informatica Cloud Service resource does not extract the targets created at run time for data sources such as Microsoft Azure Blob Storage and Azure Data Lake Store Gen2. |
| EIC-40946 | The SAP S/4HANA resource does not extract views appended to an ABAP Core Data Services (CDS) view in the catalog. |
| EIC-40459 | When you run a column profile on Teradata views, the profiling service ignores the columns with unidentified data types. |
| EIC-37244 | Null pointer exception appears during the resource connection validation for an HDFS resource in the offline scanner execution mode. The following error messages appear in the event.log file: - ERROR - java.lang.NullPointerException - ERROR - Test Connection failed. null - ERROR - Failed to test connection between resource and source |
| EIC-37225 | The Catalog Service fails to start with the following ingestion error: graphdao.JanusGraphObjectIterator: java.lang.IllegalStateException |

| Bug | Description |
|-----------|---|
| EIC-26496 | Cannot use an external PostgreSQL database for the Informatica Cluster Service. |
| EIC-25076 | You cannot view tables on the control summary page if the SQL override in the source qualifier contains a WHERE clause. |
| EIC-18754 | Permissions configured in Catalog Administrator for a user do not get applied for that user across various permission categories. |

Closed Enhancements

The following table describes closed enhancement requests:

| Issue | Description |
|-----------|---|
| EIC-41761 | Effective in version 10.4.1.1, the MicroStrategy resource creates connection parameters for a table using the database schema name. |

Enterprise Data Preparation Fixed Limitations (10.4.1.1)

The following table describes fixed limitations:

| Bug | Description |
|-----------|--|
| IDL-16902 | Cannot preview and prepare the partition files merged as a single file. |
| IDL-16815 | When you successively use the split functionality in the recipe of a worksheet, the publishing fails with the following error: <code>java.lang.OutOfMemoryError: Java heap space</code> |

Security Fixed Limitations and Closed Enhancements (10.4.1.1)

Fixed Limitations

The following table describes fixed limitations:

| Bug | Description |
|------------|--|
| PLAT-26803 | LDAP Search query filter is not correct while fetching ranged attributes during LDAP group sync. |

Closed Enhancements

The following table describes closed enhancements:

| Bug | Description |
|------------|---|
| PLAT-27008 | Enhanced support for SAML authentication using PingFederate or ADFS for Informatica web applications. |

Third-Party Fixed Limitations (10.4.1.1)

The following table describes fixed limitations:

| Bug | Description |
|------------|--|
| OCON-25799 | When you run a mapping to add columns to a Hive target table that is partitioned, the mapping fails with a validation error on Hortonworks HDP 3.x. Apache ticket reference number: HIVE-6131 |
| EIC-25198 | The DataBase Scripts resource fails with the following error while parsing the concat() function: <code>DEBUG EXPLIB_D0002 Message: 'syntax error, unexpected CONST_INTEGER_TOKEN, expecting DATATYPE_FORMAT_TOKEN or DATATYPE_SIMPLE_TOKEN or DATATYPE_TOKEN or NAMED_TOKEN</code> MITI ticket reference number: INFAEDC-1490 |
| BDM-31535 | Mappings fail intermittently in Spark mode when the Data Integration Service fails to fetch delegation token for Hive metastore. GIC case number: 02586644 |
| EIC-42359 | The SSIS resource does not extract connections for an SSIS Data Transformation Services (DTS) package. MITI ticket reference number: INFAEDC-1488 |

10.4.1.1 Known Limitations

This section contains known limitations that were found in 10.4.1.1.

Data Engineering Integration Known Limitations (10.4.1.1)

The following table describes known limitations:

| Issue | Description |
|-----------|---|
| BDM-34812 | <p>When the Spark engine runs a mapping that executes a Hive query and the tracing level is verbose, the mapping might fail with the following error:</p> <pre>SEVERE: Spark task [<task>] failed with the following error: [com.informatica.sdk.dtm.ExecutionException: [SPARK_1003] Spark task [<task>] failed with the following error: [User class threw exception: java.lang.reflect.InvocationTargetException</pre> |
| BDM-34629 | <p>Creation of a cluster configuration and Hive connection fails for a Glue-enabled Amazon EMR 5.29 cluster when the property hive-metastore-uris is not present in hive-site.xml.</p> <p>Workaround: Verify that the property hive-metastore-uris on the EMR 5.29 cluster is populated with a pointer to the Hive host. For example:</p> <pre>thrift://<Hive host name>:<port></pre> <p>Verify that the property is also present in the hive-site.xml configuration file.</p> |

Data Engineering Streaming Known Limitations (10.4.1.1)

The following table describes known limitations:

| Bug | Description |
|----------|--|
| IIS-4998 | On Cloudera CDP version 7.1.1, a streaming mapping with an Amazon Kinesis source publishes duplicate data to the target. |
| IIS-4996 | <p>When you run a streaming mapping that has a Python transformation with <code>CustomFunctionCall</code>, the mapping fails with the following error:</p> <pre>java.util.NoSuchElementException: head of empty list</pre> <p>Workaround: In the Python transformation, create either equal or more number of output ports than input ports. The downstream transformations do not need to use these output ports.</p> |
| IIS-4985 | <p>In a streaming mapping with a Kafka or Confluent Kafka source and Kafka, Confluent Kafka, or HDFS target, the mapping fails when the following conditions are true:</p> <ul style="list-style-type: none">- Target schema format is of type Flat- Mapping Flow option is enabled in the target <p>The mapping fails with the following error:</p> <pre>Error: Mapping is failed. Lost task 0.1 in stage 0.0 (TID 1,inbdsrh74qa01.informatica.com, executor 1): java.util.NoSuchElementException: next on empty iterator at scala.collection.Iterator\$\$anon\$2.next(Iterator.scala:39)</pre> |
| IIS-4973 | When you run a streaming mapping with a Kafka source and Amazon S3 target on Amazon EMR version 5.29 secure cluster, rollover of the Amazon S3 target files fail with <code>Access Denied</code> error. |

| Bug | Description |
|----------|---|
| IIS-4972 | When you run a streaming mapping with a Kafka source and Hive target on Cloudera CDP version 7.1.1, Hortonworks HDP version 3.0, and HDInsight version 4.0, any pre-task for Hive table is not accepted. Truncate table, apply new schema, and apply new column are a few examples of pre-tasks. |
| IIS-4921 | When you run a streaming mapping with a Kafka source and HDFS target on Cloudera CDP version 7.1.1, the mapping fails with a compilation error when the following conditions are true: <ul style="list-style-type: none"> - Topic name is linked to file name port and data to data port - Dynamic strategy is enabled at source - Run-time properties of the mapping is configured as <code>infaspark.hive.impersonated.complexfile.write=true</code> |
| IIS-4873 | Cleanup task fails for a streaming mapping if the mapping contains a Java transformation and a Network File System (NFS) is configured as the disTemp directory of the domain. This issue occurs on Hortonworks HDP version 3.0, HDInsight version 4.0, and Cloudera CDP version 7.1.1 clusters. Workaround: Configure the disTemp directory of the domain to point to a local file system. |

Data Privacy Management Known Limitations (10.4.1.1)

The following table describes known limitations:

| Bug | Description |
|------------|---|
| SATS-30910 | A domain discovery scan job on an Apache Cassandra data store fails if the data store includes keyspaces, tables, or columns with uppercase characters. |

Enterprise Data Catalog Known Limitations (10.4.1.1)

The following table describes known limitations:

| Bug | Description |
|-----------|---|
| EIC-46135 | While configuring a Tableau resource, if you use the search option to select a specific workbook as the repository object, and the search result fetches only one workbook, the entire project that includes the workbook is also selected. |
| EIC-46055 | Deleting an Oracle resource that includes a large volume of data assets fails with an Out Of Memory (OOM) error if Data Asset Analytics is enabled. |
| EIC-46016 | When you back up the catalog, the resource operations in progress such as resource runs and resource delete do not pause. |
| EIC-46010 | Back up of catalog using the CLI or REST API fails with the following error: Please make sure that atleast [2] vcores are available on cluster to perform this action.]. |
| EIC-45982 | If you created resources using REST API, you cannot filter resources by users in the Created By filter section of the Resource Library page. |

| Bug | Description |
|-----------|--|
| EIC-45897 | When you use AWS Glue Catalog as the Hive data source on an Amazon Elastic MapReduce version 5.29 cluster, the Hive resource scan fails with the following error: Error: org.apache.hive.service.cli.HiveSQLException: Error while processing statement: |
| EIC-43778 | When the online backup is in progress, you cannot view the asset details in the catalog. |
| EIC-43653 | Enterprise Data Catalog incorrectly displays a business term name if the business term contains a forward slash (/) in its name. This issue occurs after you assign the business term as a classified custom attribute. |
| EIC-43116 | A user gets deleted from Enterprise Data Catalog when the user follows an asset. The following error message appears in LDM.log file: ERROR[ChangeNotificationEmailThread:ChangeNotificationSender@371] - Error occurred while retrieving an email address for the user [Native\UniqueUser1] : [com.informatica.isp.corecommon.exceptions.ISPConfigException: [UM_10007] The user [UniqueUser1] in security domain [Native] |
| EIC-42898 | If you had configured SAML Single Sign-on using NetScaler and had not specified the Assertion Consumer Service URL for the applications, you can still log in to the applications. |
| DAA-1945 | You cannot view the resource scan history, user searches, and user login information from the time Data Asset Analytics is disabled to the time Data Asset Analytics is re-enabled. This issue occurs if you create and run a resource with the Save and Run option, perform searches, or log in to Enterprise Data Catalog when Data Asset Analytics is disabled. |

Profiles and Scorecards Known Limitations (10.4.1.1)

The following table describes known limitations:

| Bug | Description |
|----------|--|
| IDE-4899 | When you run a column profile on a Hive source, the profile run fails on the Spark engine if the profile has a filter and runs on a random sample. |
| IDE-4893 | When you use the infacmd ps Execute command to run a column profile in Informatica Developer, the profile run fails. |
| IDE-4887 | When you choose the Use global setting for all profiles option for an enterprise discovery profile and run it on the Spark engine, the profile run fails. |

Third-Party Known Limitations (10.4.1.1)

The following table describes known limitations:

| Bug | Description |
|------------------------|--|
| IIS-4903 | <p>When you run a streaming mapping with a Kafka source and Amazon S3 target and the Amazon S3 staging bucket name contains a dot (.), the mapping fails with the following error:</p> <pre>Unable to execute HTTP request: Certificate for xxxx doesn't match any of the subject alternative names</pre> <p>This issue occurs when you use the Cloudera CDP 7.1 distribution. Amazon S3 ticket number: 7087613701</p> |
| BDM-34608 BDM-34615 | <p>When the Spark engine runs a mapping with a Hive WASBS source and a flat file target on a non-ESP-enabled Hadoop cluster, it creates separate staging directories for users other than the expected impersonation user.</p> <p>Workaround for non-ESP-enabled HDInsight 4.0 clusters: Set the following properties on the cluster:</p> <ul style="list-style-type: none">- Set "Run as end user instead of Hive user" (doAs) to TRUE.- Enable hive.metastore.execute.setugi. <p>Workaround for ESP-enabled HDInsight 4.0 clusters: set "Run as end user instead of Hive user" (doAs) to FALSE. Microsoft ticket number: 120072723000878.</p> |
| BDM-34563 | <p>When "Create or replace schema" is enabled for a Hive table in a mapping that runs on an Amazon EMR 5.29 cluster with Glue enabled as the Hive metastore, the mapping fails with a <code>java.net.NoRouteToHostException: No route to host</code> error. The same mapping succeeds when the engine runs with a custom database and not the default database. AWS case number: 7217954611</p> |
| IDL-17211 | <p>If the partitioned column name in a parquet file is in camel case or uppercase, the data preview and preparation does not display any data. Apache case number: HIVE-23347.</p> <p>Workaround: Use lowercase for the partitioned column names.</p> |