



Informatica™

Informatica® Big Data Management
10.2.2

Big Data Management Administrator Guide

© Copyright Informatica LLC 2017, 2020

This software and documentation are provided only under a separate license agreement containing restrictions on use and disclosure. No part of this document may be reproduced or transmitted in any form, by any means (electronic, photocopying, recording or otherwise) without prior consent of Informatica LLC.

U.S. GOVERNMENT RIGHTS Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, duplication, disclosure, modification, and adaptation is subject to the restrictions and license terms set forth in the applicable Government contract, and, to the extent applicable by the terms of the Government contract, the additional rights set forth in FAR 52.227-19, Commercial Computer Software License.

Informatica and the Informatica logo are trademarks or registered trademarks of Informatica LLC in the United States and many jurisdictions throughout the world. A current list of Informatica trademarks is available on the web at <https://www.informatica.com/trademarks.html>. Other company and product names may be trade names or trademarks of their respective owners.

Portions of this software and/or documentation are subject to copyright held by third parties. Required third party notices are included with the product.

The information in this documentation is subject to change without notice. If you find any problems in this documentation, report them to us at infa_documentation@informatica.com.

Informatica products are warranted according to the terms and conditions of the agreements under which they are provided. INFORMATICA PROVIDES THE INFORMATION IN THIS DOCUMENT "AS IS" WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NON-INFRINGEMENT.

Publication Date: 2020-03-17

Table of Contents

Preface	8
Informatica Resources.	8
Informatica Network.	8
Informatica Knowledge Base.	8
Informatica Documentation.	8
Informatica Product Availability Matrices.	9
Informatica Velocity.	9
Informatica Marketplace.	9
Informatica Global Customer Support.	9
Chapter 1: Introduction to Big Data Management Administration.....	10
Big Data Management Engines.	10
Hadoop Integration.	11
Hadoop Utilities.	11
High Availability.	12
Run-time Process on the Blaze Engine.	13
Run-time Process on the Spark Engine.	14
Databricks Integration.	15
Run-time Process on the Databricks Spark Engine.	15
Application Services.	16
Data Integration Service Process.	17
Security.	18
Connections.	20
Chapter 2: Authentication.....	21
Authentication Overview.	21
Support for Authentication Systems on Hadoop.	22
Authentication with Kerberos.	23
User Impersonation.	23
Authentication with Apache Knox Gateway.	24
Chapter 3: Running Mappings on a Cluster with Kerberos Authentication	25
Running Mappings with Kerberos Authentication Overview.	25
Running Mappings in a Kerberos-Enabled Hadoop Environment.	26
Step 1. Set Up the Kerberos Configuration File on the Domain Host.	26
Step 2. Set up the Cross-Realm Trust.	28
Step 3. Create Matching Operating System Profile Names.	30
Step 4. Create the Principal Name and Keytab Files in the Active Directory Server.	30
Step 5. Specify the Kerberos Authentication Properties for the Data Integration Service.	31
Step 6. Configure the Execution Options for the Data Integration Service.	32

User Impersonation with Kerberos Authentication.	32
User Impersonation in the Hadoop Environment.	32
User Impersonation in the Native Environment.	33
Running Mappings in the Native Environment.	34
Configure the Analyst Service.	34
Chapter 4: Authorization.	35
Authorization Overview.	35
Support for Authorization Systems on Hadoop.	36
HDFS Permissions.	36
SQL Authorization for Hive.	37
Key Management Servers.	37
Configuring KMS for Informatica User Access.	37
Configuring Access to an SSL-Enabled Cluster	38
Configure the Hive Connection for SSL-Enabled Clusters.	39
Import Security Certificates from an SSL-Enabled Cluster.	39
Import Security Certificates from a TLS-Enabled Domain.	40
Chapter 5: Cluster Configuration	42
Cluster Configuration Overview.	42
Cluster Configuration and Connections.	43
Copying a Connection to Another Domain.	43
Cluster Configuration Views.	44
Active Properties View.	44
Overridden Properties View.	45
Create a Hadoop Cluster Configuration.	46
Before You Import.	46
Importing a Hadoop Cluster Configuration from the Cluster.	47
Importing a Hadoop Cluster Configuration from a File.	48
Create a Databricks Cluster Configuration.	49
Importing a Databricks Cluster Configuration from the Cluster.	49
Importing a Databricks Cluster Configuration from a File.	50
Edit the Cluster Configuration.	52
Filtering Cluster Configuration Properties.	52
Overriding Imported Properties.	53
Creating User-Defined Properties.	54
Deleting Cluster Configuration Properties.	55
Refresh the Cluster Configuration	56
Example - Cluster Configuration Refresh.	56
Delete a Cluster Configuration.	57
Cluster Configuration Privileges and Permissions.	57
Privileges and Roles.	57
Permissions.	57

Chapter 6: Cloud Provisioning Configuration	59
Cloud Provisioning Configuration Overview.	59
Verify Prerequisites	60
Enable DNS Resolution from an On-Premises Informatica Domain.	60
AWS Cloud Provisioning Configuration Properties.	61
General Properties.	61
Permissions	62
EC2 Configuration.	62
Azure Cloud Provisioning Configuration Properties.	63
Authentication Details.	63
Storage Account Details.	63
Cluster Deployment Details.	65
External Hive Metastore Details.	65
Databricks Cloud Provisioning Configuration Properties.	66
Create the Cloud Provisioning Configuration.	66
Complete the Azure Cloud Provisioning Configuration.	66
Create a Cluster Connection.	67
Chapter 7: Data Integration Service Processing.....	68
Overview of Data Integration Service Processing.	68
Data Integration Service Queueing.	69
Execution Pools.	70
Local Queues.	70
Distributed Queues.	71
Big Data Job Recovery	72
Scenarios Where Recovery is Possible.	72
Scenarios Where Recovery is Not Possible.	73
Recovery Job Management.	73
Monitoring Recovered Jobs.	73
Big Data Job Recovery Configuration.	74
Tuning for Big Data Processing.	74
Deployment Types.	75
Tuning the Application Services.	75
Tuning the Hadoop Run-time Engines.	76
Autotune.	78
Appendix A: Connections.....	81
Connections.	82
Cloud Provisioning Configuration.	82
AWS Cloud Provisioning Configuration Properties.	83
Azure Cloud Provisioning Configuration Properties.	84
Databricks Cloud Provisioning Configuration Properties.	87

Amazon Redshift Connection Properties.	88
Amazon S3 Connection Properties.	89
Cassandra Connection Properties.	91
Databricks Connection Properties.	92
Google Analytics Connection Properties.	94
Google BigQuery Connection Properties.	94
Google Cloud Spanner Connection Properties.	95
Google Cloud Storage Connection Properties.	96
Hadoop Connection Properties.	97
Hadoop Cluster Properties.	97
Common Properties.	99
Reject Directory Properties.	100
Hive Pushdown Configuration.	100
Blaze Configuration.	100
Spark Configuration.	101
HDFS Connection Properties.	102
HBase Connection Properties.	104
HBase Connection Properties for MapR-DB.	104
Hive Connection Properties.	105
JDBC Connection Properties.	108
Sqoop Connection-Level Arguments.	111
Kafka Connection Properties.	113
Microsoft Azure Blob Storage Connection Properties.	114
Microsoft Azure Cosmos DB SQL API Connection Properties.	115
Microsoft Azure Data Lake Store Connection Properties.	116
Microsoft Azure SQL Data Warehouse Connection Properties.	117
Snowflake Connection Properties.	118
Creating a Connection to Access Sources or Targets.	119
Creating a Hadoop Connection.	119
Configuring Hadoop Connection Properties.	121
Cluster Environment Variables.	121
Cluster Library Path.	121
Common Advanced Properties.	122
Blaze Engine Advanced Properties.	122
Spark Advanced Properties.	123
Appendix B: Multiple Blaze Instances on a Cluster.	127
Overview.	127
Step 1. Prepare the Hadoop Cluster for the Blaze Engine.	128
Create a Blaze User Account.	129
Create Blaze Engine Directories and Grant Permissions.	129
Grant Permissions on the Hive Source Database.	129
Step 2. Configure Data Integration Service Properties.	129

Configure Data Integration Service Process Properties.	129
Step 3. Create a Hadoop Connection.	130
Step 4. Configure Additional Hadoop Connection Properties.	131
Step 5. Set Mapping Preferences.	132
Result.	134
Appendix C: Monitoring REST API.	135
Monitoring REST API Overview.	135
Monitoring Metadata Document.	136
Sample Metadata Document.	136
MappingStats.	138
Sample Retrieve Mapping Statistics.	140
MappingAdvancedStats.	140
Sample Retrieve Advanced Mapping Statistics.	142
MappingExecutionSteps.	145
Sample Retrieve Mapping Execution Steps.	147
MappingExecutionPlans.	149
Sample Retrieve Mapping Execution Plans.	150
Index.	154

Preface

The Informatica® *Big Data Management Administrator Guide* is written for Informatica administrators. The guide contains information that you need to administer the integration of the Informatica domain with the compute clusters in non-native environments. It includes information about security, connections, and cluster configurations. This guide assumes that you are familiar with the native environment of the Informatica domain and the non-native environments where you push run-time processing.

Informatica Resources

Informatica provides you with a range of product resources through the Informatica Network and other online portals. Use the resources to get the most from your Informatica products and solutions and to learn from other Informatica users and subject matter experts.

Informatica Network

The Informatica Network is the gateway to many resources, including the Informatica Knowledge Base and Informatica Global Customer Support. To enter the Informatica Network, visit <https://network.informatica.com>.

As an Informatica Network member, you have the following options:

- Search the Knowledge Base for product resources.
- View product availability information.
- Create and review your support cases.
- Find your local Informatica User Group Network and collaborate with your peers.

Informatica Knowledge Base

Use the Informatica Knowledge Base to find product resources such as how-to articles, best practices, video tutorials, and answers to frequently asked questions.

To search the Knowledge Base, visit <https://search.informatica.com>. If you have questions, comments, or ideas about the Knowledge Base, contact the Informatica Knowledge Base team at KB_Feedback@informatica.com.

Informatica Documentation

Use the Informatica Documentation Portal to explore an extensive library of documentation for current and recent product releases. To explore the Documentation Portal, visit <https://docs.informatica.com>.

Informatica maintains documentation for many products on the Informatica Knowledge Base in addition to the Documentation Portal. If you cannot find documentation for your product or product version on the Documentation Portal, search the Knowledge Base at <https://search.informatica.com>.

If you have questions, comments, or ideas about the product documentation, contact the Informatica Documentation team at infa_documentation@informatica.com.

Informatica Product Availability Matrices

Product Availability Matrices (PAMs) indicate the versions of the operating systems, databases, and types of data sources and targets that a product release supports. You can browse the Informatica PAMs at <https://network.informatica.com/community/informatica-network/product-availability-matrices>.

Informatica Velocity

Informatica Velocity is a collection of tips and best practices developed by Informatica Professional Services and based on real-world experiences from hundreds of data management projects. Informatica Velocity represents the collective knowledge of Informatica consultants who work with organizations around the world to plan, develop, deploy, and maintain successful data management solutions.

You can find Informatica Velocity resources at <http://velocity.informatica.com>. If you have questions, comments, or ideas about Informatica Velocity, contact Informatica Professional Services at ips@informatica.com.

Informatica Marketplace

The Informatica Marketplace is a forum where you can find solutions that extend and enhance your Informatica implementations. Leverage any of the hundreds of solutions from Informatica developers and partners on the Marketplace to improve your productivity and speed up time to implementation on your projects. You can find the Informatica Marketplace at <https://marketplace.informatica.com>.

Informatica Global Customer Support

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

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

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

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

CHAPTER 1

Introduction to Big Data Management Administration

This chapter includes the following topics:

- [Big Data Management Engines, 10](#)
- [Hadoop Integration, 11](#)
- [Databricks Integration, 15](#)
- [Application Services, 16](#)
- [Security, 18](#)
- [Connections, 20](#)

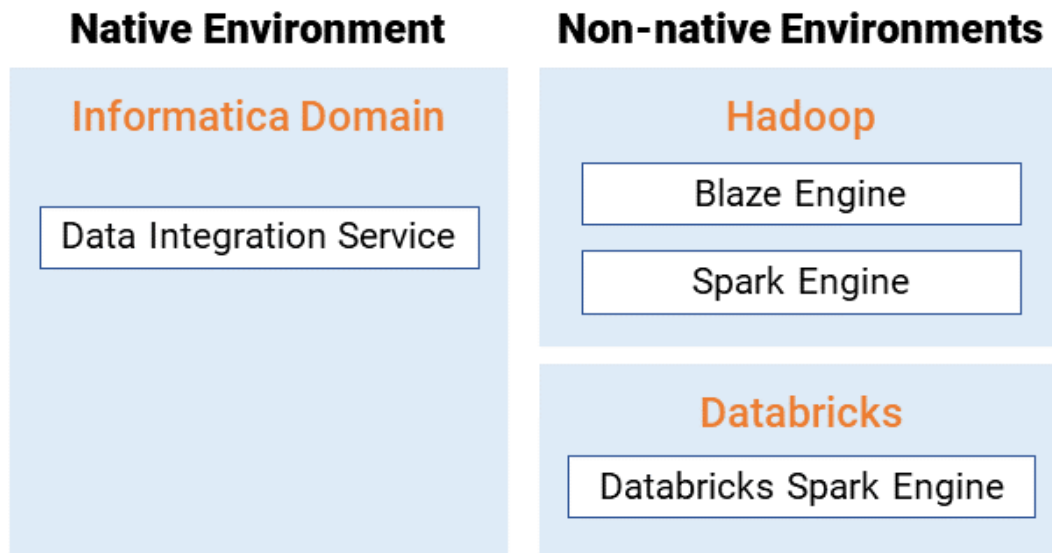
Big Data Management Engines

When you run a big data mapping, you can choose to run the mapping in the native environment or in a non-native environment, such as Hadoop or Databricks. When you validate a mapping, you can validate it against one or all of the engines. The Developer tool returns validation messages for each engine.

When you run a mapping in the native environment, the Data Integration Service in the Informatica domain runs the mapping. When you run the mapping in a non-native environment, the Data Integration Service pushes the run-time processing to a compute cluster in the non-native environment.

When you run the mapping in a non-native environment, the Data Integration Service uses a proprietary rule-based methodology to determine the best engine to run the mapping. The rule-based methodology evaluates the mapping sources and the mapping logic to determine the engine. The Data Integration Service translates the mapping logic into code that the engine can process, and it transfers the code to the engine.

The following image shows the processing environments and the run-time engines in the environments:



Hadoop Integration

Big Data Management can connect to clusters that run different Hadoop distributions. Hadoop is an open-source software framework that enables distributed processing of large data sets across clusters of machines. You might also need to use third-party software clients to set up and manage your Hadoop cluster.

Big Data Management can connect to the supported data source in the Hadoop environment, such as HDFS, HBase, or Hive, and push job processing to the Hadoop cluster. To enable high performance access to files across the cluster, you can connect to an HDFS source. You can also connect to a Hive source, which is a data warehouse that connects to HDFS.

It can also connect to NoSQL databases such as HBase, which is a database comprising key-value pairs on Hadoop that performs operations in real-time. The Data Integration Service can push mapping jobs to the Spark or Blaze engine, and it can push profile jobs to the Blaze engine in the Hadoop environment.

Big Data Management supports more than one version of some Hadoop distributions. By default, the cluster configuration wizard populates the latest supported version.

Hadoop Utilities

Big Data Management uses third-party Hadoop utilities such as Sqoop to process data efficiently.

Sqoop is a Hadoop command line program to process data between relational databases and HDFS through MapReduce programs. You can use Sqoop to import and export data. When you use Sqoop, you do not need to install the relational database client and software on any node in the Hadoop cluster.

To use Sqoop, you must configure Sqoop properties in a JDBC connection and run the mapping in the Hadoop environment. You can configure Sqoop connectivity for relational data objects, customized data

objects, and logical data objects that are based on a JDBC-compliant database. For example, you can configure Sqoop connectivity for the following databases:

- Aurora
- Greenplum
- IBM DB2
- IBM DB2 for z/OS
- Microsoft SQL Server
- Netezza
- Oracle
- Teradata
- Vertica

The Model Repository Service uses JDBC to import metadata. The Data Integration Service runs the mapping in the Hadoop run-time environment and pushes the job processing to Sqoop. Sqoop then creates map-reduce jobs in the Hadoop cluster, which perform the import and export job in parallel.

Specialized Sqoop Connectors

When you run mappings through Sqoop, you can use the following specialized connectors:

OraOop

You can use OraOop with Sqoop to optimize performance when you read data from or write data to Oracle. OraOop is a specialized Sqoop plug-in for Oracle that uses native protocols to connect to the Oracle database.

You can configure OraOop when you run Sqoop mappings on the Spark engine.

Teradata Connector for Hadoop (TDCH) Specialized Connectors for Sqoop

You can use the following TDCH specialized connectors for Sqoop to read data from or write data to Teradata:

- Cloudera Connector Powered by Teradata
- Hortonworks Connector for Teradata (powered by the Teradata Connector for Hadoop)
- MapR Connector for Teradata

These connectors are specialized Sqoop plug-ins that Cloudera, Hortonworks, and MapR provide for Teradata. They use native protocols to connect to the Teradata database.

Informatica supports Cloudera Connector Powered by Teradata and Hortonworks Connector for Teradata on the Blaze and Spark engines. When you run Sqoop mappings on the Blaze engine, you must configure these connectors. When you run Sqoop mappings on the Spark engine, the Data Integration Service invokes these connectors by default.

Informatica supports MapR Connector for Teradata on the Spark engine. When you run Sqoop mappings on the Spark engine, the Data Integration Service invokes the connector by default.

Note: For information about running native Teradata mappings with Sqoop, see the *Informatica PowerExchange for Teradata Parallel Transporter API User Guide*.

High Availability

High availability refers to the uninterrupted availability of Hadoop cluster components.

You can use high availability for the following services and security systems in the Hadoop environment on Cloudera CDH, Hortonworks HDP, and MapR Hadoop distributions:

- Apache Ranger
- Apache Ranger KMS
- Apache Sentry
- Cloudera Navigator Encrypt
- HBase
- Hive Metastore
- HiveServer2
- Name node
- Resource Manager

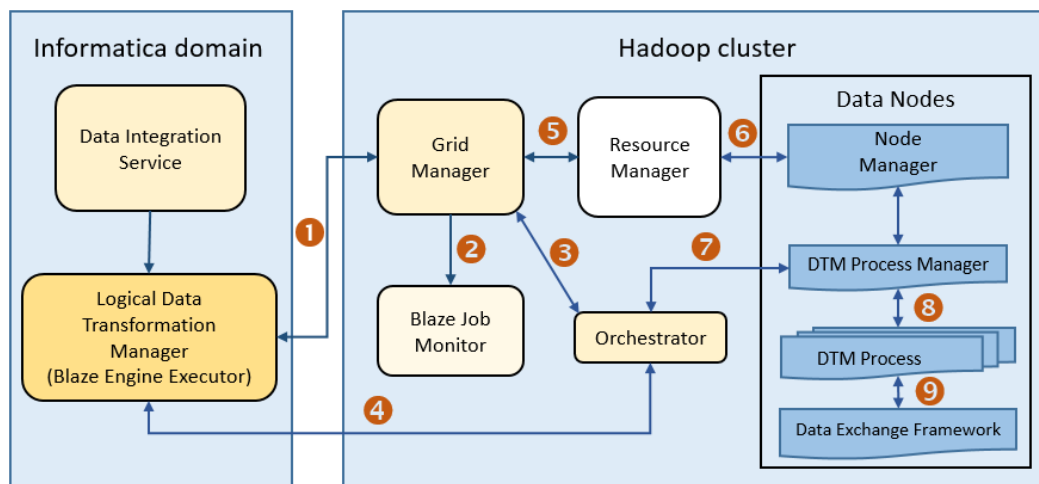
Run-time Process on the Blaze Engine

To run a mapping on the Informatica Blaze engine, the Data Integration Service submits jobs to the Blaze engine executor. The Blaze engine executor is a software component that enables communication between the Data Integration Service and the Blaze engine components on the Hadoop cluster.

The following Blaze engine components appear on the Hadoop cluster:

- Grid Manager. Manages tasks for batch processing.
- Orchestrator. Schedules and processes parallel data processing tasks on a cluster.
- Blaze Job Monitor. Monitors Blaze engine jobs on a cluster.
- DTM Process Manager. Manages the DTM Processes.
- DTM Processes. An operating system process started to run DTM instances.
- Data Exchange Framework. Shuffles data between different processes that process the data on cluster nodes.

The following image shows how a Hadoop cluster processes jobs sent from the Blaze engine executor:



The following events occur when the Data Integration Service submits jobs to the Blaze engine executor:

1. The Blaze Engine Executor communicates with the Grid Manager to initialize Blaze engine components on the Hadoop cluster, and it queries the Grid Manager for an available Orchestrator.

2. The Grid Manager starts the Blaze Job Monitor.
3. The Grid Manager starts the Orchestrator and sends Orchestrator information back to the LDTM.
4. The LDTM communicates with the Orchestrator.
5. The Grid Manager communicates with the Resource Manager for available resources for the Orchestrator.
6. The Resource Manager handles resource allocation on the data nodes through the Node Manager.
7. The Orchestrator sends the tasks to the DTM Processes through the DTM Process Manger.
8. The DTM Process Manager continually communicates with the DTM Processes.
9. The DTM Processes continually communicate with the Data Exchange Framework to send and receive data across processing units that run on the cluster nodes.

Manage Blaze Engines

The Blaze engine remains running after a mapping run. To save resources, you can set a property to stop Blaze engine infrastructure after a specified time period.

Save resources by shutting down Blaze engine infrastructure after a specified time period.

Set the `infagrid.blaze.service.idle.timeout` property or the `infagrid.orchestrator.svc.sunset.time` property. You can use the `infacmd isp createConnection` command, or set the property in the Blaze Advanced properties in the Hadoop connection in the Administrator tool or the Developer tool.

Configure the following Blaze advanced properties in the Hadoop connection:

infagrid.blaze.service.idle.timeout

Optional: The number of minutes that the Blaze engine remains idle before releasing the resources that the Blaze engine uses.

The value must be an integer. Default is 60.

infagrid.orchestrator.svc.sunset.time

Optional: Maximum lifetime for an Orchestrator service, in hours.

You can disable sunset by setting the property to 0 or a negative value. If you disable sunset, the Orchestrator never shuts down during a mapping run.

The value must be an integer. Default is 24 hours.

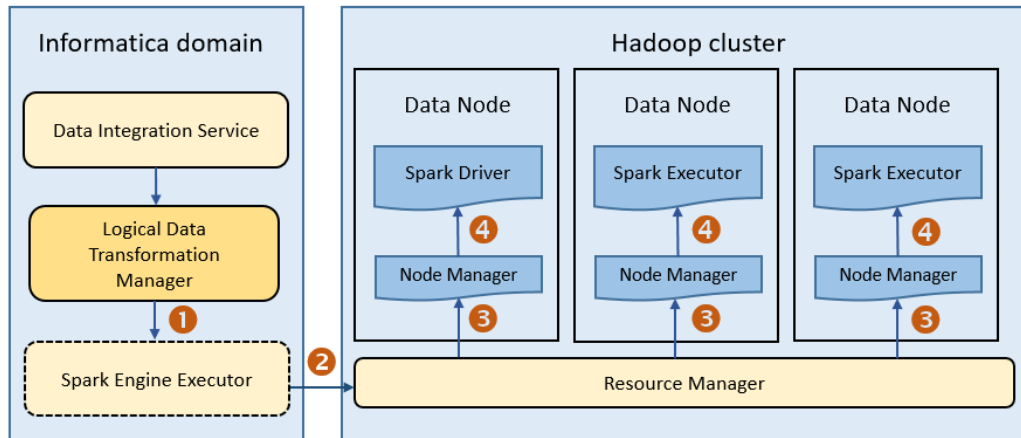
Note: Set the value to less than the Kerberos user ticket maximum renewal time. This Kerberos policy property is usually named like "Maximum lifetime for user ticket renewal."

Run-time Process on the Spark Engine

The Data Integration Service can use the Spark engine on a Hadoop cluster to run Model repository mappings.

To run a mapping on the Spark engine, the Data Integration Service sends a mapping application to the Spark executor. The Spark executor submits the job to the Hadoop cluster to run.

The following image shows how a Hadoop cluster processes jobs sent from the Spark executor:



The following events occur when Data Integration Service runs a mapping on the Spark engine:

1. The Logical Data Transformation Manager translates the mapping into a Scala program, packages it as an application, and sends it to the Spark executor.
2. The Spark executor submits the application to the Resource Manager in the Hadoop cluster and requests resources to run the application.

Note: When you run mappings on the HDInsight cluster, the Spark executor launches a spark-submit script. The script requests resources to run the application.

3. The Resource Manager identifies the Node Managers that can provide resources, and it assigns jobs to the data nodes.
4. Driver and Executor processes are launched in data nodes where the Spark application runs.

Databricks Integration

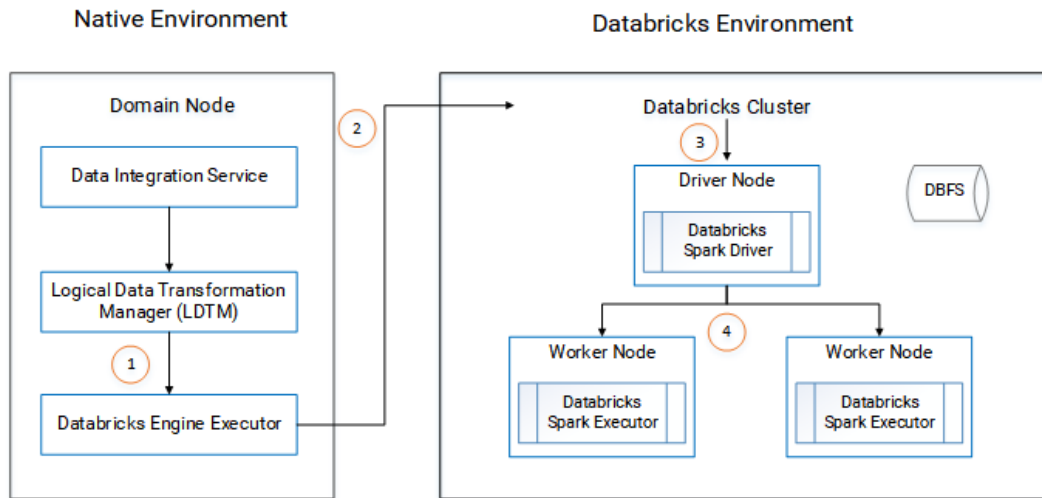
Big Data Management can push mappings to the Azure Databricks environment. Azure Databricks is an analytics cloud platform that is optimized for the Microsoft Azure cloud services. It incorporates the open-source Apache Spark cluster technologies and capabilities.

The Data Integration Service automatically installs the binaries required to integrate the Informatica domain with the Databricks environment. The integration requires Informatica connection objects and cluster configurations. A cluster configuration is a domain object that contains configuration parameters that you import from the Databricks cluster. You then associate the cluster configuration with connections to access the Databricks environment.

Run-time Process on the Databricks Spark Engine

When you run a job on the Databricks Spark engine, the Data Integration Service pushes the processing to the Databricks cluster, and the Databricks Spark engine runs the job.

The following image shows the components of the Informatica and the Databricks environments:



1. The Logical Data Transformation Manager translates the mapping into a Scala program, packages it as an application, and sends it to the Databricks Engine Executor on the Data Integration Service machine.
2. The Databricks Engine Executor submits the application through REST API to the Databricks cluster, requests to run the application, and stages files for access during run time.
3. The Databricks cluster passes the request to the Databricks Spark driver on the driver node.
4. The Databricks Spark driver distributes the job to one or more Databricks Spark executors that reside on worker nodes.
5. The executors run the job and stage run-time data to the Databricks File System (DBFS) of the workspace.

Application Services

Big Data Management uses application services in the Informatica domain to process data.

Use the Administrator tool to create connections, monitor jobs, and manage application services that Big Data Management uses.

Big Data Management uses the following application services:

Analyst Service

The Analyst Service runs the Analyst tool in the Informatica domain. The Analyst Service manages the connections between service components and the users that have access to the Analyst tool.

Data Integration Service

The Data Integration Service can process mappings in the native environment or push the mapping for processing to a compute cluster in a non-native environment. The Data Integration Service also retrieves metadata from the Model repository when you run a Developer tool mapping or workflow. The Analyst tool and Developer tool connect to the Data Integration Service to run profile jobs and store profile results in the profiling warehouse.

Mass Ingestion Service

The Mass Ingestion Service manages and validates mass ingestion specifications that you create in the Mass Ingestion tool. The Mass Ingestion Service deploys specifications to the Data Integration Service. When a specification runs, the Mass Ingestion Service generates ingestion statistics.

Metadata Access Service

The Metadata Access Service allows the Developer tool to import and preview metadata from a Hadoop cluster.

The Metadata Access Service contains information about the Service Principal Name (SPN) and keytab information if the Hadoop cluster uses Kerberos authentication. You can create one or more Metadata Access Services on a node. Based on your license, the Metadata Access Service can be highly available.

HBase, HDFS, Hive, and MapR-DB connections use the Metadata Access Service when you import an object from a Hadoop cluster. Create and configure a Metadata Access Service before you create HBase, HDFS, Hive, and MapR-DB connections.

Model Repository Service

The Model Repository Service manages the Model repository. The Model Repository Service connects to the Model repository when you run a mapping, mapping specification, profile, or workflow.

REST Operations Hub

The REST Operations Hub Service is an application service in the Informatica domain that exposes Informatica product functionality to external clients through REST APIs.

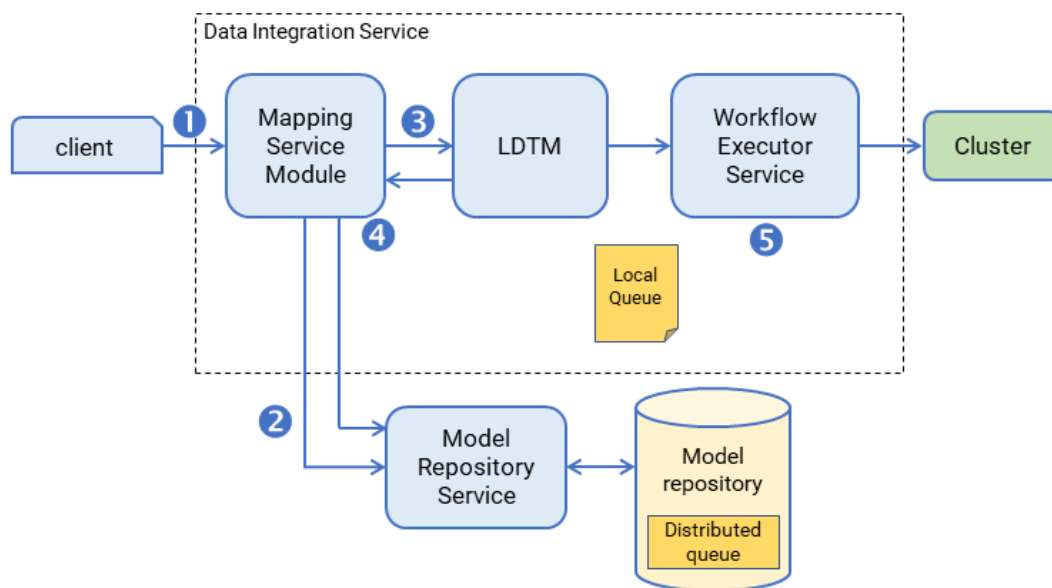
Data Integration Service Process

The Data Integration Service starts one or more Data Integration Service processes to manage requests to run mapping jobs in the Hadoop environment.

When you run mappings in the Hadoop environment, the following components run within the Data Integration Service process:

- Mapping Service Module. The Mapping Service Module receives requests to run mappings from clients.
- Logical Data Transformation Manager (LDTM). The LDTM compiles and optimizes mapping jobs, and it generates the execution workflow that is used to run a mapping on a Hadoop cluster.
- Workflow Executor Service. The Workflow Executor Service is a part of the Data Transformation Manager (DTM). The Data Integration Service uses the Workflow Executor Service to push jobs to a Hadoop cluster.

The following diagram shows how the components interact with the client, the Hadoop cluster, and the Model Repository Service:



1. A client submits a mapping execution request to the Data Integration Service. The Mapping Service Module receives the request and stores the job in the queue.
2. The Mapping Service Module connects to the Model Repository Service to fetch mapping metadata from the Model repository.
3. The Mapping Service Module passes the mapping to the Logical Data Transformation Manager (LDTM).
4. The LDTM compiles the mapping and generates the Spark execution workflow. It stores the execution workflow in the Model repository.
5. The LDTM pushes the execution workflow through the Workflow Executor Service to the cluster for processing.

For more information about the architecture of a Data Integration Service, see the "Data Integration Service Architecture" chapter in the *Informatica Application Service Guide*.

Security

Security for Big Data Management includes security for the Informatica domain native environment and for the non-native environments.

Security for the Hadoop Environment

You can configure security for Big Data Management and the Hadoop cluster to protect from threats inside and outside the network. Security for the Hadoop cluster includes the following areas:

Authentication

When the Informatica domain includes Big Data Management, user identities must be authenticated in the Informatica domain and the Hadoop cluster. Authentication for the Informatica domain is separate from authentication for the Hadoop cluster.

By default, Hadoop does not verify the identity of users. To authenticate user identities, you can configure the following authentication protocols on the cluster:

- Native authentication
- Lightweight Directory Access Protocol (LDAP)
- Kerberos, when the Hadoop distribution supports it
- Apache Knox Gateway

Big Data Management also supports Hadoop clusters that use a Microsoft Active Directory (AD) Key Distribution Center (KDC) or an MIT KDC.

Authorization

After a user is authenticated, a user must be authorized to perform actions. For example, a user must have the correct permissions to access the directories where specific data is stored to use that data in a mapping.

You can run mappings on a cluster that uses one of the following security management systems for authorization:

- Cloudera Navigator Encrypt
- HDFS permissions
- User impersonation
- Apache Ranger
- Apache Sentry
- HDFS Transparent Encryption

Data and metadata management

Data and metadata management involves managing data to track and audit data access, update metadata, and perform data lineage. Big Data Management supports Cloudera Navigator and Metadata Manager to manage metadata and perform data lineage.

Data security

Data security involves protecting sensitive data from unauthorized access. Big Data Management supports data masking with the Data Masking transformation in the Developer tool, Dynamic Data Masking, and Persistent Data Masking.

Operating system profiles

An operating system profile is a type of security that the Data Integration Service uses to run mappings. Use operating system profiles to increase security and to isolate the run-time environment for users. Big Data Management supports operating system profiles on all Hadoop distributions. In the Hadoop run-time environment, the Data Integration Service pushes the processing to the Hadoop cluster and the Big Data Management engines run mappings with the operating system profile.

Security for the Databricks Environment

The Data Integration Service uses token-based authentication to provide access to the Databricks environment. Generate tokens within the Databricks environment and use the token ID to connect to Databricks.

Connections

Big Data Management requires connections to the non-native environment.

Create and managed the following types of connections within the Administrator tool:

Cluster connection

The connection to the non-native cluster environment. You can create a Hadoop or Databricks cluster connection. The cluster connection must contain a reference either to a cloud provisioning connection or to a cluster configuration.

Cloud provisioning connection

A type of connection that enables the Data Integration Service to contact and create resources on the cloud platform. Create a cloud provisioning configuration when you configure a cluster workflow. You create a cluster workflow to automate the creation of clusters and workflow tasks on an Amazon Web Services, Microsoft Azure, or Azure Databricks cloud platform.

Cluster configuration

An object in the domain that contains configuration information about the compute cluster. The cluster configuration enables the Data Integration Service to push mapping logic to the non-native environment. Import configuration properties from the compute cluster to create a cluster configuration. You can import directly from the cluster or from an archive file. When you create the cluster configuration, you can also choose to create related connections.

CHAPTER 2

Authentication

This chapter includes the following topics:

- [Authentication Overview, 21](#)
- [Support for Authentication Systems on Hadoop, 22](#)
- [Authentication with Kerberos, 23](#)
- [Authentication with Apache Knox Gateway, 24](#)

Authentication Overview

The authentication process verifies the identity of a user account. With Big Data Management, user identities must be authenticated in the native Informatica domain and the non-native environment containing a Hadoop or Databricks cluster. Authentication for the Informatica domain is separate from authentication for the Hadoop cluster.

The Informatica domain uses native, LDAP, and Kerberos authentication. Native authentication stores user credentials and privileges in the domain configuration repository and performs all user authentication within the Informatica domain. LDAP authentication uses an LDAP directory service that stores user accounts and credentials that are accessed over the network.

Hadoop Authentication

By default, Hadoop does not authenticate users. Any user can be used in the Hadoop connection. Informatica recommends that you enable authentication for the cluster. If authentication is enabled for the cluster, the cluster authenticates the user account used for the Hadoop connection between Big Data Management and the cluster.

For a higher level of security, you can set up one of the following types of authentication for the cluster.:

Kerberos authentication

Kerberos is a network authentication protocol that uses tickets to authenticate users and services in a network. Users are stored in the Kerberos principal database, and tickets are issued by a KDC. User impersonation allows different users to run mappings on a Hadoop cluster that uses Kerberos authentication or connect to big data sources and targets that use Kerberos authentication.

Apache Knox Gateway

The Apache Knox Gateway is a REST API gateway that authenticates users and acts as a single access point for a Hadoop cluster.

For more information about how to enable authentication for the Hadoop cluster, see the documentation for your Hadoop distribution.

Databricks Authentication

The Data Integration Service uses token-based authentication to provide access to the Databricks environment. The Databricks administrator creates a token user and generates tokens for the user. The Databricks cluster configuration contains the token ID required for authentication.

Note: If the token has an expiration date, verify that you get a new token from the Databricks administrator before it expires.

Support for Authentication Systems on Hadoop

Depending on the run-time engine that you use, you can run mappings on a Hadoop cluster that uses a supported security management system.

Hadoop clusters use a variety of security management systems for user authentication. The following table shows the run-time engines supported for the security management system installed on the Hadoop platform:

Hadoop Distribution	Apache Knox	Kerberos	LDAP
Amazon EMR	No support	- Native - Blaze - Spark	- Native - Blaze - Spark
Azure HDInsight	No support	- Native - Blaze - Spark	No support
Cloudera CDH	No support	- Native - Blaze - Spark	- Native - Blaze - Spark
Hortonworks HDP	- Native - Blaze - Spark	- Native - Blaze - Spark	- Native - Blaze - Spark
MapR	No support	- Native - Blaze - Spark	No support

Additional Information for Authentication

Consider the following additional security support statements:

- Informatica supports an Azure HDInsight cluster with Enterprise Security Package. The Enterprise Security Package uses Kerberos for authentication and Apache Ranger for authorization.
- Sqoop cannot access Kerberos-enabled databases.

Authentication with Kerberos

Big Data Management and the Hadoop cluster can use Kerberos authentication to verify user accounts, when the Hadoop cluster supports Kerberos. You can use Kerberos authentication with the Informatica domain, with a supported Hadoop cluster, or with both.

Kerberos is a network authentication protocol that uses tickets to authenticate access to services and nodes in a network. Kerberos uses a Key Distribution Center (KDC) to validate the identities of users and services and to grant tickets to authenticated user and service accounts. Users and services are known as principals. The KDC has a database of principals and their associated secret keys that are used as proof of identity. Kerberos can use an LDAP directory service as a principal database.

You can integrate the Informatica domain with a Kerberos-enabled Hadoop cluster whether the domain is Kerberos-enabled or not.

The requirements for Kerberos authentication for the Informatica domain and for the Hadoop cluster:

Kerberos authentication for the Informatica domain

Kerberos authentication for the Informatica domain requires principals stored in a Microsoft Active Directory (AD) LDAP service. If the Informatica domain is Kerberos-enabled, you must use Microsoft AD for the KDC.

Kerberos authentication for the Hadoop cluster

Informatica supports Hadoop clusters that use an AD KDC or an MIT KDC.

When you enable Kerberos for Hadoop, each user and Hadoop service must be authenticated by the KDC. The cluster must authenticate the Data Integration Service user and, optionally, the Blaze user.

For more information about how to configure Kerberos for Hadoop, see the documentation for your Hadoop distribution.

The configuration steps required for Big Data Management to connect to a Hadoop cluster that uses Kerberos authentication depend on whether the Informatica domain uses Kerberos.

User Impersonation

User impersonation allows different users to run mappings in a Hadoop cluster that uses Kerberos authentication or connect to big data sources and targets that use Kerberos authentication.

The Data Integration Service uses its credentials to impersonate the user accounts designated in the Hadoop connection to connect to the Hadoop cluster or to start the Blaze engine.

When the Data Integration Service impersonates a user account to submit a mapping, the mapping can only access Hadoop resources that the impersonated user has permissions on. Without user impersonation, the Data Integration Service uses its credentials to submit a mapping to the Hadoop cluster. Restricted Hadoop resources might be accessible.

When the Data Integration service impersonates a user account to start the Blaze engine, the Blaze engine has the privileges and permissions of the user account used to start it.

Authentication with Apache Knox Gateway

The Apache Knox Gateway is a REST API gateway that authenticates users and acts as a single access point for a Hadoop cluster.

Knox creates a perimeter around a Hadoop cluster. Without Knox, users and applications must connect directly to a resource in the cluster, which requires configuration on the client machines. A direct connection to resources exposes host names and ports to all users and applications and decreases the security of the cluster.

If the cluster uses Knox, applications use REST APIs and JDBC/ODBC over HTTP to connect to Knox. Knox authenticates the user and connects to a resource.

CHAPTER 3

Running Mappings on a Cluster with Kerberos Authentication

This chapter includes the following topics:

- [Running Mappings with Kerberos Authentication Overview, 25](#)
- [Running Mappings in a Kerberos-Enabled Hadoop Environment, 26](#)
- [User Impersonation with Kerberos Authentication, 32](#)
- [Running Mappings in the Native Environment, 34](#)
- [Configure the Analyst Service, 34](#)

Running Mappings with Kerberos Authentication Overview

You can run mappings on a Hadoop cluster that uses MIT or Microsoft Active Directory (AD) Kerberos authentication. Kerberos is a network authentication protocol that uses tickets to authenticate access to services and nodes in a network.

If the Informatica domain uses Kerberos authentication, you must configure a one-way cross-realm trust to enable the Hadoop cluster to communicate with the Informatica domain. The Informatica domain uses Kerberos authentication on an AD service. The Hadoop cluster uses Kerberos authentication on an MIT service. Enable the cross-realm trust to enable the MIT service to communicate with the AD service.

Based on whether the Informatica domain uses Kerberos authentication or not, you might need to perform the following tasks to run mappings on a Hadoop cluster that uses Kerberos authentication:

- If you run mappings in a Hadoop environment, you must configure user impersonation to enable other users to run mappings on the Hadoop cluster.
- If you run mappings in the native environment, you must configure the mappings to read and process data from Hive sources that use Kerberos authentication.
- If you run a mapping that has Hive sources or targets, you must enable user authentication for the mapping on the Hadoop cluster.
- If you import metadata from Hive, complex file sources, and HBase sources, you must configure the Developer tool to use Kerberos credentials to access the Hive, complex file, and HBase metadata.

Running Mappings in a Kerberos-Enabled Hadoop Environment

To run mappings in a Kerberos-enabled Hadoop environment, you must configure the Kerberos configuration file, create user authentication artifacts, and configure Kerberos authentication properties for the Informatica domain.

The Kerberos configuration file `krb5.conf` contains configuration properties for the Kerberos realm. The one-way cross-realm trust enables the Informatica domain to communicate with the Hadoop cluster.

The Informatica domain uses Kerberos authentication on a Microsoft Active Directory service. The Hadoop cluster uses Kerberos authentication on an MIT Kerberos service. You set up a one-way cross-realm trust to enable the KDC for the MIT Kerberos service to communicate with the KDC for the Active Directory service. After you set up the cross-realm trust, you must configure the Informatica domain to enable mappings to run in the Hadoop cluster.

To run mappings on a cluster that uses Kerberos authentication, perform the following configuration tasks:

1. Set up the Kerberos configuration file.
2. When the Informatica domain uses Kerberos authentication, set up the one-way cross-realm trust.
3. Create matching operating system profile user names on each Hadoop cluster node.
4. Create the Service Principal Name and Keytab File in the Active Directory Server.
5. Specify the Kerberos authentication properties for the Data Integration Service.
6. Configure Execution Options for the Data Integration Service.

Step 1. Set Up the Kerberos Configuration File on the Domain Host

Set the configuration properties for the Kerberos realm that the Hadoop cluster uses to `krb5.conf` on the machine on which the Data Integration Service runs.

`krb5.conf` is located in the `<Informatica Installation Directory>/java/jre/lib/security` directory.

1. Back up `krb5.conf` before you make any changes.
2. Open `krb5.conf` for editing.
3. In the `libdefaults` section, set the following properties:
 - `default_realm`. Name of the service realm for the Informatica domain. The value is the same whether or not the domain uses Kerberos authentication.
 - `udp_preference_limit`. Determines the protocol that Kerberos uses when it sends a message to the KDC. Set to 1 to use the TCP protocol.

The following example shows the value if the Informatica domain does not use Kerberos authentication:

```
[libdefaults]
default_realm = hadoop-realm.example.com
udp_preference_limit=1
```

The following example shows the value if the Informatica domain uses Kerberos authentication:

```
[libdefaults]
default_realm = INFA-AD-REALM.example.com
udp_preference_limit=1
```

4. In the `realms` section, set or add the properties required by Informatica.

The following table lists the values to which you must set properties in the realms section:

Parameter	Value
kdc	Name of the host running a KDC server for that realm.
admin_server	Name of the Kerberos administration server.

The following example shows the parameters for the Hadoop realm if the Informatica domain does not use Kerberos authentication:

```
[realms]
HADOOP-REALM = {
    kdc = 123abcd134.hadoop-realm.com
    admin server = def456.hadoop-realm.com
}
```

The following example shows the parameters for the Hadoop realm if the Informatica domain uses Kerberos authentication:

```
[realms]
INFA-AD-REALM = {
    kdc = 123abcd.infa-realm.com
    admin server = 123abcd.infa-realm.com
}
HADOOP-REALM = {
    kdc = 123abcd134.hadoop-realm.com
    admin server = def456.hadoop-realm.com
}
```

5. In the *domain_realms* section, map the domain name or host name to a Kerberos realm name. The domain name is prefixed by a period (.).

The following example shows the parameters for the Hadoop domain_realm if the Informatica domain does not use Kerberos authentication:

```
[domain_realm]
.hadoop_realm.com = HADOOP-REALM
hadoop_realm.com = HADOOP-REALM
```

The following example shows the parameters for the Hadoop domain_realm if the Informatica domain uses Kerberos authentication:

```
[domain_realm]
.infa_ad_realm.com = INFA-AD-REALM
infa_ad_realm.com = INFA-AD-REALM
.hadoop_realm.com = HADOOP-REALM
hadoop_realm.com = HADOOP-REALM
```

6. Copy the `krb5.conf` file to the following locations on the machine that hosts the Data Integration Service:

- <Informatica installation directory>/services/shared/security/
- <Informatica installation directory>/java/jre/lib/security

The following example shows the content of `krb5.conf` with the required properties for an Informatica domain that does not use Kerberos authentications:

```
[libdefaults]
default_realm = HADOOP-REALM
udp_preference_limit=1

[realms]
HADOOP-REALM = {
    kdc = 123abcd134.hadoop-realm.com
    admin_server = 123abcd124.hadoop-realm.com
}
```

```
[domain_realm]
.hadoop_realm.com = HADOOP-REALM
hadoop_realm.com = HADOOP-REALM
```

The following example shows the content of `krb5.conf` with the required properties for an Informatica domain that uses Kerberos authentication:

```
[libdefaults]
default_realm = INFA-AD-REALM
udp_preference_limit=1

[realms]
INFA-AD-REALM = {
    kdc = abc123.infa-ad-realm.com
    admin_server = abc123.infa-ad-realm.com
}

HADOOP-REALM = {
    kdc = def456.hadoop-realm.com
    admin_server = def456.hadoop-realm.com
}

[domain_realm]
.infa_ad_realm.com = INFA-AD-REALM
infa_ad_realm.com = INFA-AD-REALM
.hadoop_realm.com = HADOOP-REALM
hadoop_realm.com = HADOOP-REALM
```

Step 2. Set up the Cross-Realm Trust

Perform this step when the Informatica domain uses Kerberos authentication.

Set up a one-way cross-realm trust to enable the KDC for the MIT Kerberos server to communicate with the KDC for the Active Directory server. When you set up the one-way cross-realm trust, the Hadoop cluster can authenticate the Active Directory principals.

To set up the cross-realm trust, you must complete the following steps:

1. Configure the Active Directory server to add the local MIT realm trust.
2. Configure the MIT server to add the cross-realm principal.
3. Translate principal names from the Active Directory realm to the MIT realm.

Configure the Microsoft Active Directory Server

Add the MIT KDC host name and local realm trust to the Active Directory server.

To configure the Active Directory server, complete the following steps:

1. Enter the following command to add the MIT KDC host name:

```
ksetup /addkdc <mit_realm_name> <kdc_hostname>
```

For example, enter the command to add the following values:

```
ksetup /addkdc HADOOP-MIT-REALM def456.hadoop-mit-realm.com
```

2. Enter the following command to add the local realm trust to Active Directory:

```
netdom trust <mit_realm_name> /Domain:<ad_realm_name> /add /realm /
passwordt:<TrustPassword>
```

For example, enter the command to add the following values:

```
netdom trust HADOOP-MIT-REALM /Domain:INFA-AD-REALM /add /realm /passwordt:trust1234
```

3. Enter the following commands based on your Microsoft Windows environment to set the proper encryption type:

For Microsoft Windows 2008, enter the following command:

```
ksetup /SetEncTypeAttr <mit_realm_name> <enc_type>
```

For Microsoft Windows 2003, enter the following command:

```
ktpass /MITRealmName <mit_realm_name> /TrustEncryp <enc_type>
```

Note: The `enc_type` parameter specifies AES, DES, or RC4 encryption. To find the value for `enc_type`, see the documentation for your version of Windows Active Directory. The encryption type you specify must be supported on both versions of Windows that use Active Directory and the MIT server.

Configure the MIT Server

Configure the MIT server to add the cross-realm `krbtgt` principal. The `krbtgt` principal is the principal name that a Kerberos KDC uses for a Windows domain.

Enter the following command in the `kadmin.local` or `kadmin` shell to add the cross-realm `krbtgt` principal:

```
kadmin: addprinc -e "<enc_type_list>" krbtgt/<mit_realm_name>@<MY-AD-REALM.COM>
```

The `enc_type_list` parameter specifies the types of encryption that this cross-realm `krbtgt` principal will support. The `krbtgt` principal can support either AES, DES, or RC4 encryption. You can specify multiple encryption types. However, at least one of the encryption types must correspond to the encryption type found in the tickets granted by the KDC in the remote realm.

For example, enter the following value:

```
kadmin: addprinc -e "rc4-hmac:normal des3-hmac-shal:normal" krbtgt/HADOOP-MIT-  
REALM@INFA-AD-REALM
```

Translate Principal Names from the Active Directory Realm to the MIT Realm

To translate the principal names from the Active Directory realm into local names within the Hadoop cluster, you must configure the `hadoop.security.auth_to_local` property in the `core-site.xml` file and `hadoop.kms.authentication.kerberos.name.rules` property in the `kms-site.xml` file on all the machines in the Hadoop cluster.

For example, set the following property in `core-site.xml` on all the machines in the Hadoop cluster:

```
<property>  
  <name>hadoop.security.auth_to_local</name>  
  <value>  
    RULE: [1:$1@$0] (^.*@INFA-AD-REALM$) s/^ (.*)@INFA-AD-REALM$/$1/g  
    RULE: [2:$1@$0] (^.*@INFA-AD-REALM$) s/^ (.*)@INFA-AD-REALM$/$1/g  
    DEFAULT  
  </value>  
</property>
```

For example, set the following property in `kms-site.xml` on all the machines in the Hadoop cluster:

```
<property>  
  <name>hadoop.kms.authentication.kerberos.name.rules</name>  
  <value>  
    RULE: [1:$1@$0] (^.*@INFA-AD-REALM$) s/^ (.*)@INFA-AD-REALM$/$1/g  
    RULE: [2:$1@$0] (^.*@INFA-AD-REALM$) s/^ (.*)@INFA-AD-REALM$/$1/g  
    DEFAULT  
  </value>  
</property>
```

Step 3. Create Matching Operating System Profile Names

Create matching operating system profile user names on the machine that runs the Data Integration Service and each Hadoop cluster node to run Informatica mapping jobs.

For example, if user joe runs the Data Integration Service on a machine, you must create the user joe with the same operating system profile on each Hadoop cluster node.

Open a UNIX shell and enter the following UNIX command to create a user with the user name joe.

Step 4. Create the Principal Name and Keytab Files in the Active Directory Server

Create an SPN in the KDC database for Microsoft Active Directory service that matches the user name of the user that runs the Data Integration Service. Create a keytab file for the SPN on the machine on which the KDC server runs. Then, copy the keytab file to the machine on which the Data Integration Service runs.

You do not need to use the Informatica Kerberos SPN Format Generator to generate a list of SPNs and keytab file names. You can create your own SPN and keytab file name.

To create an SPN and Keytab file in the Active Directory server, complete the following steps:

Create a user in the Microsoft Active Directory Service.

Login to the machine on which the Microsoft Active Directory Service runs and create a user with the same name as the user you created in ["Step 3. Create Matching Operating System Profile Names" on page 30](#).

Create an SPN associated with the user.

Use the following guidelines when you create the SPN and keytab files:

- The user principal name (UPN) must be the same as the SPN.
- Enable delegation in Microsoft Active Directory.
- Use the ktpass utility to create an SPN associated with the user and generate the keytab file.

For example, enter the following command:

```
ktpass -out infa_hadoop.keytab -mapuser joe -pass tempBG@2008 -princ joe/
domain12345@INFA-AD-REALM -crypto all
```

Note: The `-out` parameter specifies the name and path of the keytab file. The `-mapuser` parameter is the user to which the SPN is associated. The `-pass` parameter is the password for the SPN in the generated keytab. The `-princ` parameter is the SPN.

- Use the ktutil utility to generate the keytab file for an Azure HDInsight cluster that uses Enterprise Security Package and ADLS storage.

For example, enter the following command:

```
sshuser@hn0-hivesc:/tmp/keytabs$ ktutil
ktutil: addent -password -p alice -k 1 -e RC4-HMAC
Password for alice@SECUREHADOOPRC.ONMICROSOFT.COM;
ktutil: wkt /tmp/keytabs/alice.keytab
ktutil: q
```

Step 5. Specify the Kerberos Authentication Properties for the Data Integration Service

In the Data Integration Service properties, configure the properties that enable the Data Integration Service to connect to a Hadoop cluster that uses Kerberos authentication. Use the Administrator tool to set the Data Integration Service properties.

Property	Description
Cluster Staging Directory	The directory on the cluster where the Data Integration Service pushes the binaries to integrate the native and non-native environments and to store temporary files during processing. Default is /tmp.
Hadoop Staging User	The HDFS user that performs operations on the Hadoop staging directory. The user requires write permissions on Hadoop staging directory. Default is the operating system user that starts the Informatica daemon.
Custom Hadoop OS Path	<p>The local path to the Informatica server binaries compatible with the Hadoop operating system. Required when the Hadoop cluster and the Data Integration Service are on different supported operating systems. The Data Integration Service uses the binaries in this directory to integrate the domain with the Hadoop cluster. The Data Integration Service can synchronize the following operating systems:</p> <ul style="list-style-type: none"> - SUSE and Redhat <p>Include the source directory in the path. For example, <Informatica server binaries>/source.</p> <p>Changes take effect after you recycle the Data Integration Service.</p> <p>Note: When you install an Informatica EBF, you must also install it in this directory.</p>
Hadoop Kerberos Service Principal Name	<p>Service Principal Name (SPN) of the Data Integration Service to connect to a Hadoop cluster that uses Kerberos authentication.</p> <p>Not required for the MapR distribution.</p>
Hadoop Kerberos Keytab	<p>The file path to the Kerberos keytab file on the machine on which the Data Integration Service runs.</p> <p>Not required for the MapR distribution.</p>
Custom Properties	<p>Properties that are unique to specific environments.</p> <p>You can configure run-time properties for the Hadoop environment in the Data Integration Service, the Hadoop connection, and in the mapping. You can override a property configured at a high level by setting the value at a lower level. For example, if you configure a property in the Data Integration Service custom properties, you can override it in the Hadoop connection or in the mapping. The Data Integration Service processes property overrides based on the following priorities:</p> <ol style="list-style-type: none"> 1. Mapping custom properties set using infacmd ms runMapping with the -cp option 2. Mapping run-time properties for the Hadoop environment 3. Hadoop connection advanced properties for run-time engines 4. Hadoop connection advanced general properties, environment variables, and classpaths 5. Data Integration Service custom properties

Step 6. Configure the Execution Options for the Data Integration Service

To determine whether the Data Integration Service runs jobs in separate operating system processes or in one operating system process, configure the Launch Job Options property. Use the Administrator tool to configure the execution options for the Data Integration Service.

1. Click **Edit** to edit the **Launch Job Options** property in the execution options for the Data Integration Service properties.
2. Choose the launch job option.
 - If you configure the Data Integration Service to launch jobs as a separate process, you must specify the location of the `krb5.conf` file in the Java Virtual Manager (JVM) Options as a custom property in the Data Integration Service process. `krb5.conf` is located in the following directory:<Informatica Installation Directory>/java/jre/lib/security.
 - If you configure the Data Integration Service to launch jobs in the service process, you must specify the location of `krb5.conf` in the **Java Command Line Options** property in the Advanced Properties of the Data Integration Service process. Use the following syntax:

```
-Djava.security.krb5.conf=<Informatica installation directory>/java/jre/lib/security/krb5.conf
```

User Impersonation with Kerberos Authentication

You can enable different users to run mappings in a Hadoop cluster that uses Kerberos authentication or connect to big data sources and targets that use Kerberos authentication. To enable different users to run mappings or connect to big data sources and targets, you must configure user impersonation.

You can configure user impersonation for the native or Hadoop environment.

Before you configure user impersonation, you must complete the following prerequisites:

- Complete the tasks for running mappings in a Kerberos-enabled Hadoop environment.
- Configure Kerberos authentication for the native or Hadoop environment.
- If the Hadoop cluster uses MapR, create a proxy directory for the user who will impersonate other users.

If the Hadoop cluster does not use Kerberos authentication, you can specify a user name in the Hadoop connection to enable the Data Integration Service to impersonate that user.

If the Hadoop cluster uses Kerberos authentication, you must specify a user name in the Hadoop connection.

User Impersonation in the Hadoop Environment

To enable different users to run mapping and workflow jobs on a Hadoop cluster that uses Kerberos authentication, you must configure user impersonation in the Hadoop environment.

For example, you want to enable user Bob to run mappings and workflows on the Hadoop cluster that uses Kerberos authentication.

To enable user impersonation, you must complete the following steps:

1. In the Active Directory, enable delegation for the Service Principal Name for the Data Integration Service to enable Bob to run Hadoop jobs.

2. If the service principal name (SPN) is different from the impersonation user, grant read permission on Hive tables to the SPN user.
3. Specify Bob as the user name in the Hadoop connection.

User Impersonation in the Native Environment

To enable different users to run mappings that read or processes data from big data sources or targets that use Kerberos authentication, configure user impersonation for the native environment.

To enable user impersonation, you must complete the following steps:

1. Specify Kerberos authentication properties for the Data Integration Service.
2. Configure the execution options for the Data Integration Service.

Step 1. Specify the Kerberos Authentication Properties for the Data Integration Service

In the Data Integration Service properties, configure the properties that enable the Data Integration Service to connect to a Hadoop cluster that uses Kerberos authentication. Use the Administrator tool to set the Data Integration Service properties.

Description	Property
Hadoop Kerberos Service Principal Name	Service Principal Name (SPN) of the Data Integration Service to connect to a Hadoop cluster that uses Kerberos authentication. Not required for the MapR distribution.
Hadoop Kerberos Keytab	The file path to the Kerberos keytab file on the machine on which the Data Integration Service runs. Not required for the MapR distribution.

Step 2. Configure the Execution Options for the Data Integration Service

To determine whether the Data Integration Service runs jobs in separate operating system processes or in one operating system process, configure the **Launch Job Options** property. Use the Administrator tool to configure the execution options for the Data Integration Service.

1. Click **Edit** to edit the **Launch Job Options** property in the execution options for the Data Integration Service properties.
2. Choose the launch job option.
 - If you configure the Data Integration Service to launch jobs as a separate process, you must specify the location of the `krb5.conf` in the Java Virtual Manager (JVM) Options as a custom property in the Data Integration Service process. `krb5.conf` is located in the following directory: `<Informatica Installation Directory>/java/jre/lib/security`.
 - If you configure the Data Integration Service to launch jobs in the service process, you must specify the location of `krb5.conf` in the **Java Command Line Options** property in the Advanced Properties of the Data Integration Service process. Use the following syntax:

```
-Djava.security.krb5.conf=<Informatica installation directory>/java/jre/lib/security/krb5.conf
```

Running Mappings in the Native Environment

To read and process data from Hive, HBase, or HDFS sources that use Kerberos authentication, you must configure Kerberos authentication for mappings in the native environment.

To read and process data from Hive, HBase, or HDFS sources, perform the following steps:

1. Complete the tasks for running mappings in a Kerberos-enabled Hadoop environment.
2. Complete the tasks for running mappings in the Hadoop environment when Informatica uses Kerberos authentication.
3. Create matching operating system profile user names on the machine that runs the Data Integration Service and each Hadoop cluster node used to run Informatica mapping jobs.
4. Create an Active Directory user that matches the operating system profile user you created in step 3.
5. Create an SPN associated with the user.

Use the following guidelines when you create the SPN and keytab files:

- The UPN must be the same as the SPN.
- Enable delegation in Active Directory.
- Use the `ktpass` utility to create an SPN associated with the user and generate the keytabs file.

For example, enter the following command:

```
ktpass -out infa_hadoop.keytab -mapuser joe -pass tempBG@2008 -princ joe/  
domain12345@HADOOP-AD-REALM -crypto all
```

The `-out` parameter specifies the name and path of the keytab file. The `-mapuser` parameter is the user to which the SPN is associated. The `-pass` parameter is the password for the SPN in the generated keytab. The `-princ` parameter is the SPN.

Configure the Analyst Service

To use the Analyst Service with a Hadoop cluster that uses Kerberos authentication, configure the Analyst Service to use the Kerberos ticket that the Data Integration Service uses.

In the Administrator tool, select the Analyst Service. In the **Processes** tab, edit the Advanced Properties to add the following value to the JVM Command Line Options field: `DINFA_HADOOP_DIST_DIR=<Informatica installation directory>/services/shared/hadoop/<hadoop_distribution>`.

CHAPTER 4

Authorization

This chapter includes the following topics:

- [Authorization Overview, 35](#)
- [Support for Authorization Systems on Hadoop, 36](#)
- [HDFS Permissions, 36](#)
- [SQL Authorization for Hive, 37](#)
- [Key Management Servers, 37](#)
- [Configuring Access to an SSL-Enabled Cluster , 38](#)

Authorization Overview

Authorization controls what a user can do on a Hadoop cluster. For example, a user must be authorized to submit jobs to the Hadoop cluster.

You can use the following systems to manage authorization for Big Data Management:

HDFS permissions

By default, Hadoop uses HDFS permissions to determine what a user can do to a file or directory on HDFS. Additionally, Hadoop implements transparent data encryption in HDFS directories.

Apache Sentry

Sentry is a security plug-in that you can use to enforce role-based authorization for data and metadata on a Hadoop cluster. You can enable high availability for Sentry in the Hadoop cluster. Sentry can secure data and metadata at the table and column level. For example, Sentry can restrict access to columns that contain sensitive data and prevent unauthorized users from accessing the data.

Apache Ranger

Ranger is a security plug-in that you can use to authenticate users of a Hadoop cluster. Ranger manages access to files, folders, databases, tables, and columns. When you perform an action, Ranger verifies that the user meets the policy requirements and has the correct permissions on HDFS. You can enable high availability for Ranger in the Hadoop cluster.

Fine Grained Authorization

Fine grained authorization enables database administrators to impose column-level authorization on Hive tables and views. A more fine-grained level of authorization enables administrators to impose row and column level authorization. You can configure a Hive connection to observe fine grained authorization.

Support for Authorization Systems on Hadoop

Depending on the run-time engine that you use, you can run mappings on a Hadoop cluster that uses a supported security management system.

Hadoop clusters use a variety of security management systems for user authorization. The following table shows the run-time engines supported for the security management system installed on the Hadoop platform:

Hadoop Distribution	Apache Ranger	Apache Sentry	HDFS Transparent Encryption	SSL/TLS	SQL Authorization
Amazon EMR	No support	No support	No support	No support	No support
Azure HDInsight	- Native - Blaze - Spark	No support	No support	No support	- Native - Blaze - Spark
Cloudera CDH	No support	- Native - Blaze - Spark	- Native - Blaze - Spark	- Native - Blaze - Spark	- Native - Blaze - Spark
Hortonworks HDP	- Native - Blaze - Spark	No support	- Native - Blaze - Spark	- Native - Blaze - Spark	- Native - Blaze - Spark
MapR	No support	No support	No support	- Native - Blaze - Spark	No support

Additional Information for Authorization

Consider the following additional security support statements:

- Informatica supports an Azure HDInsight cluster with Enterprise Security Package. The Enterprise Security Package uses Kerberos for authentication and Apache Ranger for authorization.
- Sqoop cannot access SSL-enabled databases.
- The combination of Apache Ranger and SQL authorization is supported on Hortonworks HDP only.
- The combination of Apache Sentry and SQL authorization is supported on Cloudera, RedHat and SUSE, only.

HDFS Permissions

HDFS permissions determine what a user can do to files and directories stored in HDFS. To access a file or directory, a user must have permission or belong to a group that has permission.

HDFS permissions are similar to permissions for UNIX or Linux systems. For example, a user requires the *r* permission to read a file and the *w* permission to write a file.

When a user or application attempts to perform an action, HDFS checks if the user has permission or belongs to a group with permission to perform that action on a specific file or directory.

SQL Authorization for Hive

SQL standards-based authorization enables database administrators to impose fine grained authorization on Hive tables and views when you read data from a Hive source or a target.

Informatica supports fine grained authorization for Hive sources with Blaze engine, and Hive sources and targets with Spark engines. You can use the Ranger authorization plug-in when you enable fine grained authorization for mappings that run on a Hortonworks HDP cluster.

You can use the Sentry authorization plug-in when you enable fine grained authorization for mappings that run on a Cloudera cluster. When the mapping accesses Hive sources in Blaze engine and Hive sources and targets in Spark engine on a cluster that uses Sentry authorization and runs in native mode, you can use SQL authorization on the column level if you configure `hive.server2.proxy.user` in the Hive JDBC connect string.

In this case, the mapping uses the `hive.server2.proxy.user` value to access Hive sources and targets. When you also configure the `mappingImpersonationUserName` property, then the mapping uses the `mappingImpersonationUserName` value to access Hive sources and targets.

You can configure a Hive connection to observe fine grained authorization.

Key Management Servers

Key Management Server (KMS) is an open source key management service that supports HDFS data at rest encryption. You can use the cluster administration utility to configure the KMS for Informatica user access.

You can use the following key management servers to encrypt the data at rest:

- Apache Ranger KMS. Ranger Key Management Store is an open source, scalable cryptographic key management service that supports HDFS data at rest encryption.
- Cloudera Java KMS. For Cloudera CDH clusters, Cloudera provides a Key Management Server based on the Hadoop KeyProvider API to support HDFS data at rest encryption.
- Cloudera Navigator Encrypt. Cloudera Navigator Encrypt is a Cloudera proprietary key management service that secures the data and implements HDFS data at rest encryption.

KMS enables the following functions:

Key management

You can create, update, or delete encryption key zones that control access to functionality.

Access control policies

You can administer access control policies for encryption keys. You can create or edit keys to control access by users to functionality.

Configuring KMS for Informatica User Access

If you use a KMS to encrypt HDFS data at rest, use the cluster administration utility to configure the KMS for Informatica user access.

1. Create a KMS user account for the Informatica user. Add the Informatica user to a new KMS repository, or to an existing KMS repository.

The user corresponds to the Data Integration Service user or the Kerberos SPN user.

2. Grant permissions to the Informatica user.
3. Create and configure an encryption key.
4. Create an encryption zone that uses the encryption key you created.

For example:

```
hdfs dfs -mkdir /zone_encr_infa
hdfs crypto -createZone -keyName infa_key -path /zone_encr_infa
```

5. Browse to the Custom KMS Site page and add the following properties:

```
hadoop.kms.proxyuser.<user>.groups=*
hadoop.kms.proxyuser.<user>.hosts=*
hadoop.kms.proxyuser.<user>.users=*
```

where <user> is the Informatica user name you configured in Step 1.

6. Update the following properties:

```
hadoop.kms.proxyuser.<user>.hosts
hadoop.kms.proxyuser.<user>.groups
```

7. Search for *proxyuser* in the KMS Configurations area. To register all Hadoop system users with the KMS, add the following properties:

```
hadoop.kms.proxyuser.HTTP.hosts=*
hadoop.kms.proxyuser.HTTP.users=*
hadoop.kms.proxyuser.hive.hosts=*
hadoop.kms.proxyuser.hive.users=*
hadoop.kms.proxyuser.keyadmin.hosts=*
hadoop.kms.proxyuser.keyadmin.users=*
hadoop.kms.proxyuser.nn.hosts=*
hadoop.kms.proxyuser.nn.users=*
hadoop.kms.proxyuser.rm.hosts=*
hadoop.kms.proxyuser.rm.users=*
hadoop.kms.proxyuser.yarn.hosts=*
hadoop.kms.proxyuser.yarn.users=*
```

Configuring Access to an SSL-Enabled Cluster

When you use an SSL-enabled cluster, you must configure the Informatica domain to communicate with the secure cluster.

Based on the cluster distribution that uses SSL, you perform the following tasks:

Cloudera CDH or Hortonworks HDP cluster uses SSL

Import security certificates from the cluster to the Informatica domain. If you created a Hive connection object manually, configure the connection string properties to access the SSL-enabled cluster.

MapR cluster uses SSL

Make sure that the MapR client is configured to communicate with a secure cluster. If you created a Hive connection object manually, configure the connection string properties to access the SSL-enabled cluster.

Configure the Hive Connection for SSL-Enabled Clusters

If you created the Hive connection when you created cluster configurations, the cluster configuration creation wizard enables access to a cluster that uses SSL. If you manually created a Hive connection, you must configure the connection string properties to enable access to a cluster that uses SSL.

If you manually created a Hive connection, add the following property-value pair to the metadata connection string and data access connection string properties:

```
ssl=true
```

For example:

```
jdbc:hive2://<hostname>:<port>/<db>;ssl=true
```

Note: Insert the `ssl=true` flag before the kerberos principal element when you create the Hive connection manually.

Import Security Certificates from an SSL-Enabled Cluster

When you use custom, special, or self-signed security certificates to secure the Hadoop cluster, Informatica services that connect to the cluster require these certificates to be present on the machines that run the application services. Use the `keytool` utility to import certificates from the cluster.

For more information about the `keytool` utility, refer to the Oracle documentation.

Note: If a MapR cluster is SSL-enabled, you do not have to import the security certificates. Make sure that the MapR client on the Data Integration Service and Metadata Access Service machines is configured to access an SSL-enabled cluster. For more information about installing and configuring the MapR client, see the *Informatica Big Data Management Hadoop Integration Guide*.

If a Cloudera CDH or Hortonworks HDP cluster uses SSL, import security certificates from the cluster to the Data Integration Service and Metadata Access Service machines.

1. Run the following `keytool -exportcert` command on the cluster to export the certificates:

```
keytool -exportcert
  -alias <alias name>
  -keystore <custom.truststore file location>
  -file <exported certificate file location>
  -storepass <password>
```

Where:

- `-alias` specifies the alias name associated with the truststore file.
- `-keystore` specifies the location of the truststore file on the cluster.
- `-file` specifies the file name and location for the exported certificate file.
- `-storepass` specifies the password for the keystore on the cluster.

The `keytool -exportcert` command produces a certificate file associated with the alias.

2. Run the following `keytool -importcert` command on the Data Integration Service and Metadata Access Service machines to import the security certificates:

```
keytool -importcert -trustcacerts
  -alias <alias name>
  -file <exported certificate file location>
  -keystore <java cacerts location>
  -storepass <password>
```

Where:

- `-alias` specifies the alias name associated with the certificate file.

- -file specifies the file name and location of the exported certificate file.
- -keystore specifies the location of the truststore file on the domain.
- -storepass specifies the password for the keystore on the domain.

Depending on whether the Informatica domain uses SSL, you specify the keystore location as follows:

- If the domain is SSL-enabled, import the certificate file to the following location:
`<Informatica installation directory>/services/shared/security/infa_truststore.jks`
- If the domain is not SSL-enabled, import the certificate file to the following location:
`<Informatica installation directory>/java/jre/lib/security/cacerts`

The `keytool -importcert` command imports the security certificates to the keystore location you specify.

Example. Import Security Certificates

The big data environment includes a Cloudera CDH cluster that uses SSL and an Informatica domain that does not use SSL. You export the security certificate for the user `bigdata_user1` from the `custom.keystore` on the Cloudera CDH cluster to the file `exported.cer`. Then, you import the `export.cer` certificate file to the Informatica domain location.

1. Run the following export command:

```
keytool -exportcert -alias bigdata_user1 -keystore ~/custom.truststore -file ~/exported.cer
```

2. Run the following import command on the Data Integration Service machine:

```
keytool -importcert -alias bigdata_user1 -file ~/exported.cer -keystore <Informatica installation directory>/java/jre/lib/security/cacerts
```

3. Run the following import command on the Metadata Access Service machine:

```
keytool -importcert -alias bigdata_user1 -file ~/exported.cer -keystore <Informatica installation directory>/java/jre/lib/security/cacerts
```

Import Security Certificates from a TLS-Enabled Domain

When an Azure HDInsight cluster uses ADLS storage and the domain is configured to use TLS, you must import the certificates to the default or custom truststore file that the Informatica domain uses.

Default truststore file

If the domain is TLS-enabled and the Azure HDInsight cluster that uses ADLS as a storage uses server managed keys, you must import the Baltimore CyberTrust Root certificate to the default truststore file.

Use the `keytool` utility to import the security certificate.

The default truststore file is located in the following directory: `<Informatica installation home>/services/shared/security/infa_truststore.jks`

Custom truststore file

If the domain is TLS-enabled and the Azure HDInsight cluster that uses ADLS as a storage uses server managed keys, get the custom truststore file location from Informatica Administrator, and then import the Baltimore CyberTrust Root certificate to the custom truststore file.

Use the `keytool` utility to import the security certificate.

To get the custom truststore file location, perform the following steps:

1. In the Administrator tool, click the Manage tab.
2. Click the Services and Nodes view.
3. In the Domain Navigator, select the domain.

4. Get the custom truststore file location from the domain properties.

You can download the Baltimore CyberTrust Root certificates from <https://www.digicert.com/digicert-root-certificates.htm>.

For more information about downloading the certificates, see <https://docs.microsoft.com/en-us/azure/java-add-certificate-ca-store>.

CHAPTER 5

Cluster Configuration

This chapter includes the following topics:

- [Cluster Configuration Overview, 42](#)
- [Cluster Configuration and Connections, 43](#)
- [Cluster Configuration Views, 44](#)
- [Create a Hadoop Cluster Configuration, 46](#)
- [Create a Databricks Cluster Configuration, 49](#)
- [Edit the Cluster Configuration, 52](#)
- [Refresh the Cluster Configuration , 56](#)
- [Delete a Cluster Configuration, 57](#)
- [Cluster Configuration Privileges and Permissions, 57](#)

Cluster Configuration Overview

A cluster configuration is an object in the domain that contains configuration information about the compute cluster. The cluster configuration enables the Data Integration Service to push mapping logic to the non-native environment. Import configuration properties from the compute cluster to create a cluster configuration. You can import directly from the cluster or from an archive file.

When you create a cluster configuration associated with the Hadoop environment, the **Cluster Configuration** wizard can create Hadoop, HBase, HDFS, and Hive connections to access the Hadoop environment. When you create a cluster configuration associated with the Databricks environment, the **Cluster Configuration** wizard can create the Databricks connection. If you choose to create the connections, the wizard also associates the configuration object with the connections.

You can override the property values, and you can create user-defined properties based on your requirements. When property values change on the cluster, you can refresh the cluster configuration, either directly from the cluster or from an archive file.

The cluster configuration contains properties for the cluster type and version. You can edit the version property to any supported version. After you change the version, you must restart the Data Integration Service.

Consider the following high-level process to manage cluster configurations:

1. Import the cluster configuration, choosing to create connections that require the cluster configuration.
2. Edit the cluster configuration. Override imported property values and add user-defined properties.

3. Refresh the cluster configuration. When property values change on the cluster, refresh the cluster configuration to import the changes.

Cluster Configuration and Connections

When you create a cluster configuration, you can choose to create connections. All cluster type connections used to run a mapping must be associated with a cluster configuration.

If you choose to create connections, the **Cluster Configuration** wizard associates the cluster configuration with each connection that it creates.

The wizard creates the following connections when you create a cluster configuration associated with the Hadoop environment:

- Hive
- HBase
- Hadoop
- HDFS

The wizard creates a Databricks connection when you create a cluster configuration associated with the Databricks environment.

The wizard uses the following naming convention when it creates connections: <connection type>_<cluster configuration name>, such as Hive_ccMapR.

If you do not choose to create connections, you must manually create them and associate the cluster configuration with them.

Copying a Connection to Another Domain

When you copy a connection associated with a cluster configuration to another domain, you must first create a cluster configuration in the target domain.

Create a cluster configuration with the same name as the one that is associated with the connection in the source domain.

Note: When you create a cluster configuration in the domain, use the **Cluster Configuration** wizard. Informatica does not recommend importing the archive file into the domain to create a cluster configuration. If you import the archive file into the domain, the user-defined properties are converted to imported properties. When you subsequently refresh the cluster configuration, the refresh operation replaces the values of properties with cluster property values, and removes properties that do not exist on the cluster.

1. Identify a connection to copy, and note the name of the cluster configuration that is associated with it.
2. In the target domain, create a cluster configuration of the same name.
3. Choose not to create connections with creation of the cluster configuration.
4. Copy the connection to the target domain.

The connection has an associated cluster configuration.

Cluster Configuration Views

You can manage cluster configurations on the **Connections** tab of the Administrator tool.

When you highlight a cluster configuration on the **Connections** tab, you can view the cluster configuration details on the right pane. The cluster configuration displays the following components:

Active Properties view

Displays general properties and all run-time properties and values. You can view the following properties:

- Imported properties with imported values and any overridden values.
- User-defined properties and values. User-defined property values appear as overridden values.

Overridden Properties view

Displays only properties with overridden values, including imported properties and user-defined properties.

Permissions view

Configure permissions to perform actions on cluster configurations.

Actions menu

From the Actions menu, you can perform the following actions:

- Refresh the cluster configuration from the Hadoop cluster.
- Export the configuration to an archive file, required by the Developer tool to access cluster metadata at design-time.

You can create, edit, and delete properties from the **Active Properties** view and the **Overridden Properties** view.

Note: The **Active Properties** and **Overridden Properties** views display configuration sets with names based on the associated *-site.xml file on the cluster. For example, the properties from the cluster core-site.xml file appear under the configuration set name core_site_xml.

Active Properties View

The **Active Properties** view displays all cluster properties, both imported and user-defined.

The **Active Properties** view contains general properties and all run-time properties. General properties include the cluster configuration name, ID, description, distribution type, and the last date of refresh. Run-time properties are organized into configuration sets based on the corresponding *-site.xml files on the cluster. For example, the hive-site.xml configuration set contains all of the properties and values imported from the hive-site.xml file on the cluster.

The cluster configuration can contain the following types of run-time properties:

Imported properties

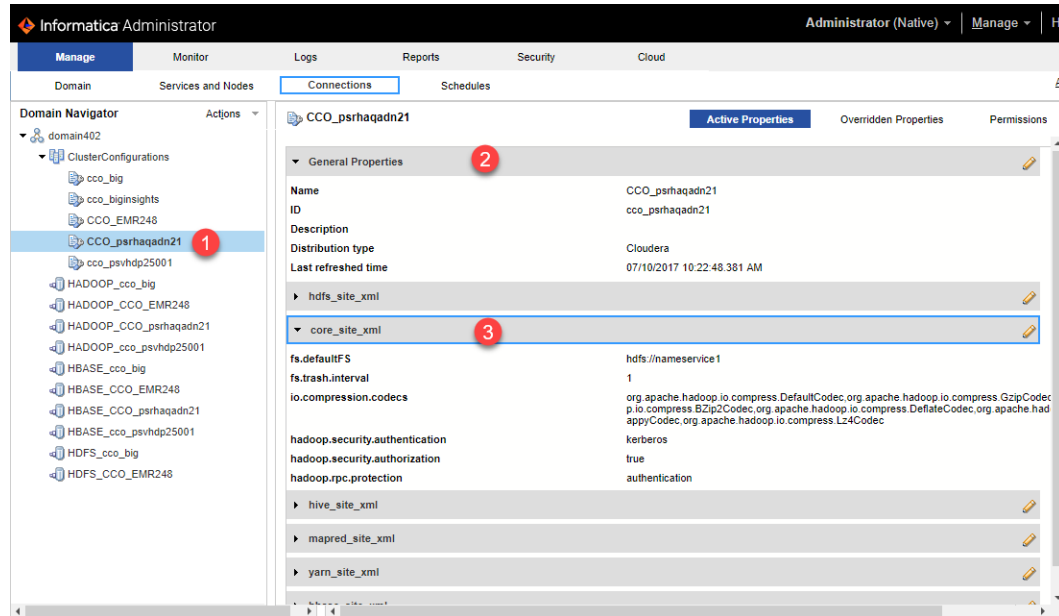
Properties and values imported from the cluster or file. You can override property values based on your requirements. Some cluster configuration properties contain sensitive information, such as passwords. The Service Manager masks the value of sensitive properties with asterisk characters when you import or refresh the cluster configuration. The masked values appear in the Administrator tool and in infacmd results.

User-defined properties

You can create user-defined properties based on processing requirements. When you create a user-defined property, the value appears as an overridden value.

Active properties are properties that the Data Integration Service uses at run time. Each expanded configuration set of the **Active Properties** view displays these active values. If a property has an overridden value, the Data Integration Service uses the overridden value as the active value. If the property does not have an overridden value, the Data Integration Service uses the imported value as the active value. To see the imported value of a property that is overridden, click the edit icon.

The following image shows cluster configurations in the **Domain Navigator**.



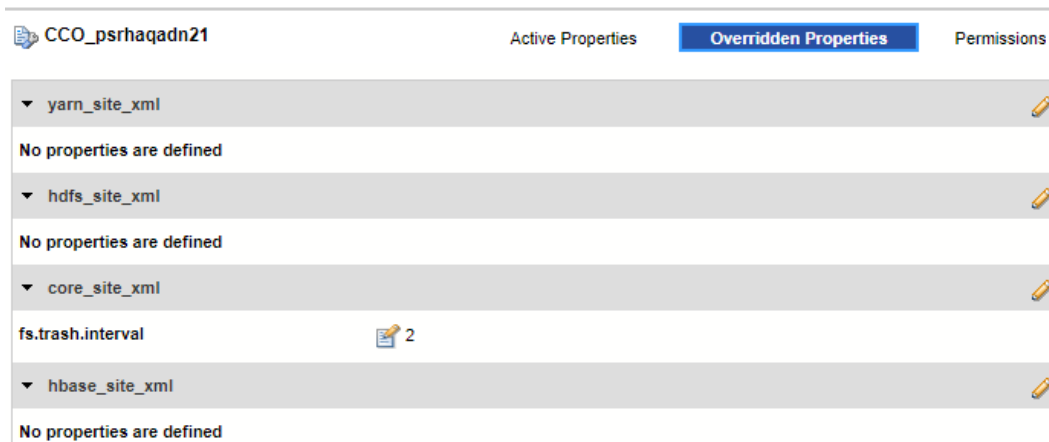
1. The **Cluster Configurations** node in the **Domain Navigator** displays the cluster configurations in the domain.
2. The right pane shows the general properties and configuration sets. The **General Properties** set is expanded to show general property values.
3. The **core-site.xml** configuration set is expanded to show the properties that it contains.

Overridden Properties View

The **Overridden Properties** view displays only properties with overridden values.

The **Overridden Properties** view includes user-defined properties and imported properties that you overrode. The values that appear in the view are the active values. To see imported values, click the edit icon.

The following image shows a property in the core-site.xml configuration set with an overridden value of 2:



Note that configuration sets that do not contain overrides display a message indicating that no properties are defined.

Create a Hadoop Cluster Configuration

Import the Hadoop cluster information into the domain. When you import cluster information, you import values from `*-site.xml` files to create a domain object called a cluster configuration.

Choose one of the following options to import cluster properties:

Import from cluster

When you import directly from the cluster, you enter cluster connection information. The Service Manager uses the information to connect to the cluster and get cluster configuration properties.

Note: You can import directly from Azure HDInsight, Cloudera CDH, and Hortonworks HDP clusters.

Import from file

When you import from a file, you browse to an archive file that the Hadoop administrator created. Use this option if the Hadoop administrator requires you to do so.

Note: If you import from a MapR or Amazon EMR cluster, you must import from a file.

Before You Import

Before you can import the cluster configuration, you must get information from the Hadoop administrator, based on the method of import.

If you import directly from the cluster, contact the Hadoop administrator to get cluster connection information. If you import from a file, get an archive file of exported cluster information.

For more information about required cluster information, see the *Big Data Management Integration Guide*.

Note: To import from Amazon EMR or MapR, you must import from an archive file.

Importing a Hadoop Cluster Configuration from the Cluster

When you import the Hadoop cluster configuration directly from the cluster, you provide information to connect to the cluster.

Get cluster connection information from the Hadoop administrator.

1. From the **Connections** tab, click the **ClusterConfigurations** node in the Domain Navigator.
2. From the Actions menu, select **New > Cluster Configuration**.

The **Cluster Configuration** wizard opens.

3. Configure the following General properties:

Property	Description
Cluster configuration name	Name of the cluster configuration.
Description	Optional description of the cluster configuration.
Distribution type	The cluster Hadoop distribution type.
Distribution version	<p>Version of the Hadoop distribution.</p> <p>Each distribution type has a default version. The default version is the latest version of the Hadoop distribution that Big Data Management supports.</p> <p>Note: When the cluster version differs from the default version and Informatica supports more than one version, the cluster configuration import process populates the property with the most recent supported version. For example, consider the case where Informatica supports versions 5.10 and 5.13, and the cluster version is 5.12. In this case, the cluster configuration import process populates this property with 5.10, because 5.10 is the most recent supported version before 5.12.</p> <p>You can edit the property to choose any supported version. Restart the Data Integration Service for the changes to take effect.</p>
Method to import the cluster configuration	Choose Import from cluster .
Create connections	<p>Choose to create Hadoop, HDFS, Hive, and HBase connections.</p> <p>If you choose to create connections, the Cluster Configuration wizard associates the cluster configuration with each connection that it creates.</p> <p>The Hadoop connection contains default values for properties such as cluster environment variables, cluster path variables, and advanced properties. Based on the cluster environment and the functionality that you use, you can add to the default values or change the default values of these properties. For a list of Hadoop connection properties to configure, see "Configuring Hadoop Connection Properties" on page 121.</p> <p>If you do not choose to create connections, you must manually create them and associate the cluster configuration with them.</p> <p>Important: When the wizard creates the Hive connection, it populates the Metadata Connection String and the Data Access Connection String properties with the value from the hive.metastore.uris property. If the Hive metastore and HiveServer2 are running on different nodes, you must update the Metadata Connection String to point to the HiveServer2 host.</p>

The cluster properties appear.

- Configure the following properties:

Property	Description
Host	IP address of the cluster manager.
Port	Port of the cluster manager.
User ID	Cluster user ID.
Password	Password for the user.
Cluster name	Name of the cluster. Use the display name if the cluster manager manages multiple clusters. If you do not provide a cluster name, the wizard imports information based on the default cluster.

- Click **Next** and verify the cluster configuration information on the summary page.

Importing a Hadoop Cluster Configuration from a File

You can import properties from an archive file to create a cluster configuration.

Before you import from the cluster, you must get the archive file from the Hadoop administrator.

- From the **Connections** tab, click the **ClusterConfigurations** node in the Domain Navigator.
- From the Actions menu, select **New > Cluster Configuration**.

The **Cluster Configuration** wizard opens.

- Configure the following properties:

Property	Description
Cluster configuration name	Name of the cluster configuration.
Description	Optional description of the cluster configuration.
Distribution type	The cluster Hadoop distribution type.
Distribution version	<p>Version of the Hadoop distribution.</p> <p>Each distribution type has a default version. This is the latest version of the Hadoop distribution that Big Data Management supports.</p> <p>When the cluster version differs from the default version, the cluster configuration wizard populates the cluster configuration Hadoop distribution property with the most recent supported version relative to the cluster version. For example, suppose Informatica supports versions 5.10 and 5.13, and the cluster version is 5.12. In this case, the wizard populates the version with 5.10.</p> <p>You can edit the property to choose any supported version. Restart the Data Integration Service for the changes to take effect.</p>

Property	Description
Method to import the cluster configuration	Choose Import from file to import properties from an archive file.
Create connections	<p>Choose to create Hadoop, HDFS, Hive, and HBase connections.</p> <p>If you choose to create connections, the Cluster Configuration wizard associates the cluster configuration with each connection that it creates.</p> <p>The Hadoop connection contains default values for properties such as cluster environment variables, cluster path variables, and advanced properties. Based on the cluster environment and the functionality that you use, you can add to the default values or change the default values of these properties. For a list of Hadoop connection properties to configure, see "Configuring Hadoop Connection Properties" on page 121.</p> <p>If you do not choose to create connections, you must manually create them and associate the cluster configuration with them.</p> <p>Important: When the wizard creates the Hive connection, it populates the Metadata Connection String and the Data Access Connection String properties with the value from the hive.metastore.uris property. If the Hive metastore and HiveServer2 are running on different nodes, you must update the Metadata Connection String to point to the HiveServer2 host.</p>

4. Click **Browse** to select a file. Select the file and click **Open**.
5. Click **Next** and verify the cluster configuration information on the summary page.

Create a Databricks Cluster Configuration

A Databricks cluster configuration is an object in the domain that contains configuration information about the Databricks cluster. The cluster configuration enables the Data Integration Service to push mapping logic to the Databricks environment.

Use the Administrator tool to import configuration properties from the Databricks cluster to create a cluster configuration. You can import configuration properties from the cluster or from a file that contains cluster properties. You can choose to create a Databricks connection when you perform the import.

Importing a Databricks Cluster Configuration from the Cluster

When you import the cluster configuration directly from the cluster, you provide information to connect to the cluster.

Before you import the cluster configuration, get cluster information from the Databricks administrator.

1. From the **Connections** tab, click the **ClusterConfigurations** node in the Domain Navigator.
2. From the Actions menu, select **New > Cluster Configuration**.
The **Cluster Configuration** wizard opens.

- Configure the following properties:

Property	Description
Cluster configuration name	Name of the cluster configuration.
Description	Optional description of the cluster configuration.
Distribution type	The distribution type. Choose Databricks .
Method to import the cluster configuration	Choose Import from cluster .
Databricks domain	Domain name of the Databricks deployment.
Databricks access token	The token ID created within Databricks required for authentication. Note: If the token has an expiration date, verify that you get a new token from the Databricks administrator before it expires.
Databricks cluster ID	The cluster ID of the Databricks cluster.
Create connection	Choose to create a Databricks connection. If you choose to create a connection, the Cluster Configuration wizard associates the cluster configuration with the Databricks connection. If you do not choose to create a connection, you must manually create one and associate the cluster configuration with it.

- Click **Next** to verify the information on the summary page.

Importing a Databricks Cluster Configuration from a File

You can import properties from an archive file to create a cluster configuration.

Complete the following tasks to import a Databricks cluster from a file:

- Get required cluster properties from the Databricks administrator.
- Create an .xml file with the cluster properties, and compress it into a .zip or .tar file.
- Log in to the Administrator tool and import the file.

Create the Import File

To import the cluster configuration from a file, you must create an archive file.

To create the .xml file for import, you must get required information from the Databricks administrator. You can provide any name for the file and store it locally.

The following table describes the properties required to import the cluster information:

Property Name	Description
cluster_name	Name of the Databricks cluster.
cluster_ID	The cluster ID of the Databricks cluster.

Property Name	Description
baseUrl	URL to access the Databricks cluster.
accesstoken	The token ID created within Databricks required for authentication.

Optionally, you can include other properties specific to the Databricks environment.

When you complete the .xml file, compress it into a .zip or .tar file for import.

Sample Import File

The following text shows a sample import file with the required properties:

```
<?xml version="1.0" encoding="UTF-8"?><configuration>
  <property>
    <name>cluster_name</name>
    <value>my_cluster</value>
  </property>
  <property>
    <name>cluster_id</name>
    <value>0926-294544-bckt123</value>
  </property>
  <property>
    <name>baseUrl</name>
    <value>https://provide.adatabricks.net/</value>
  </property>
  <property>
    <name>accesstoken</name>
    <value>dapicf76c2d4567c6sldn654fe875936e778</value>
  </property>
</configuration>
```

Import the Cluster Configuration

After you create the .xml file with the cluster properties, use the Administrator tool to import into the domain and create the cluster configuration.

1. From the **Connections** tab, click the **ClusterConfigurations** node in the Domain Navigator.
2. From the Actions menu, select **New > Cluster Configuration**.
The **Cluster Configuration** wizard opens.
3. Configure the following properties:

Property	Description
Cluster configuration name	Name of the cluster configuration.
Description	Optional description of the cluster configuration.
Distribution type	The distribution type. Choose Databricks .
Method to import the cluster configuration	Choose Import from file .

Property	Description
Upload configuration archive file	The full path and file name of the file. Click the Browse button to navigate to the file.
Create connection	Choose to create a Databricks connection. If you choose to create a connection, the Cluster Configuration wizard associates the cluster configuration with the Databricks connection. If you do not choose to create a connection, you must manually create one and associate the cluster configuration with it.

- Click **Next** to verify the information on the summary page.

Edit the Cluster Configuration

You can edit property values in a cluster configuration. You can also add user-defined properties, override imported property values, and delete properties within a configuration set.

To edit the cluster configuration, you can access the **Edit** dialog box for a configuration set on the **Active Properties** view or the **Overridden Properties** view. Within the **Edit** dialog box, you can use the filter control to find properties.

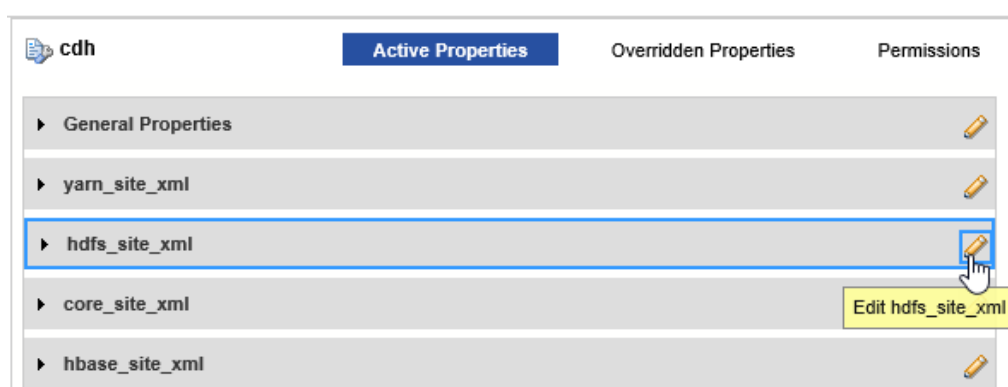
If you edit properties in a cluster configuration that the Data Integration Service used to run a mapping, recycle the Data Integration Service for changes to take effect. For example, if you change the Hadoop distribution version for a cluster configuration that the Data Integration Service used to run a mapping, recycle the Data Integration Service for the change to take effect.

Filtering Cluster Configuration Properties

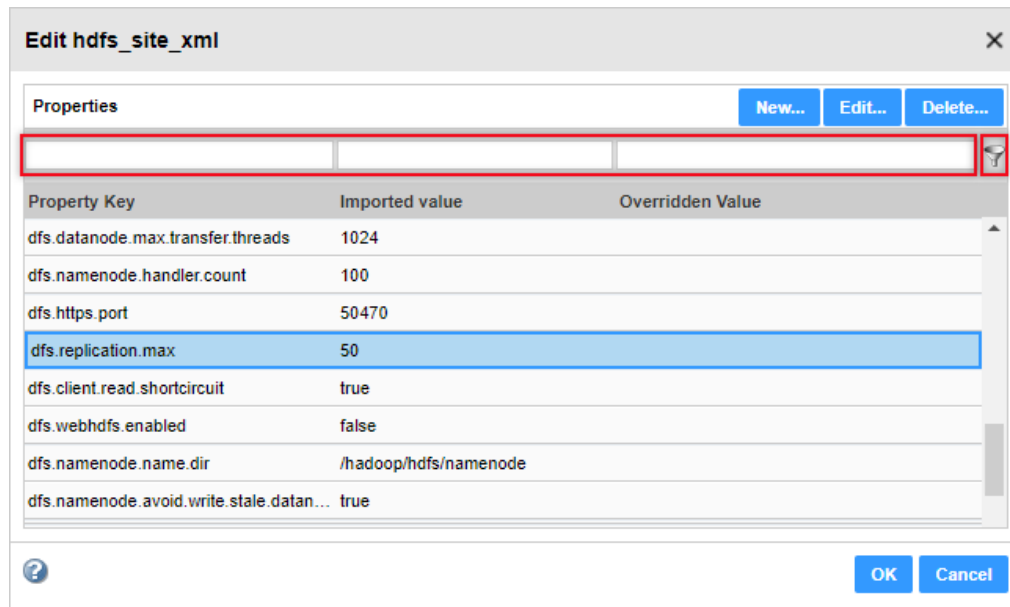
You can search for properties within a configuration set by using the filter controls.

You can filter properties in the **Active Properties** view or the **Overridden Properties** view. You might want to filter properties when a configuration set contains a large number of properties.

- In the **Active Properties** view or the **Overridden Properties** view, expand a configuration set.
- Click the **Edit** icon on the name bar of the configuration set that you want to edit.
The following image shows the **Edit** icon for the `hdfs-site.xml` configuration set:



3. Enter text in the filter text entry pane above any column, and then click the filter icon. You can search by property, imported value, overridden value.
The following image shows the filter text entry panes and the filter icon:



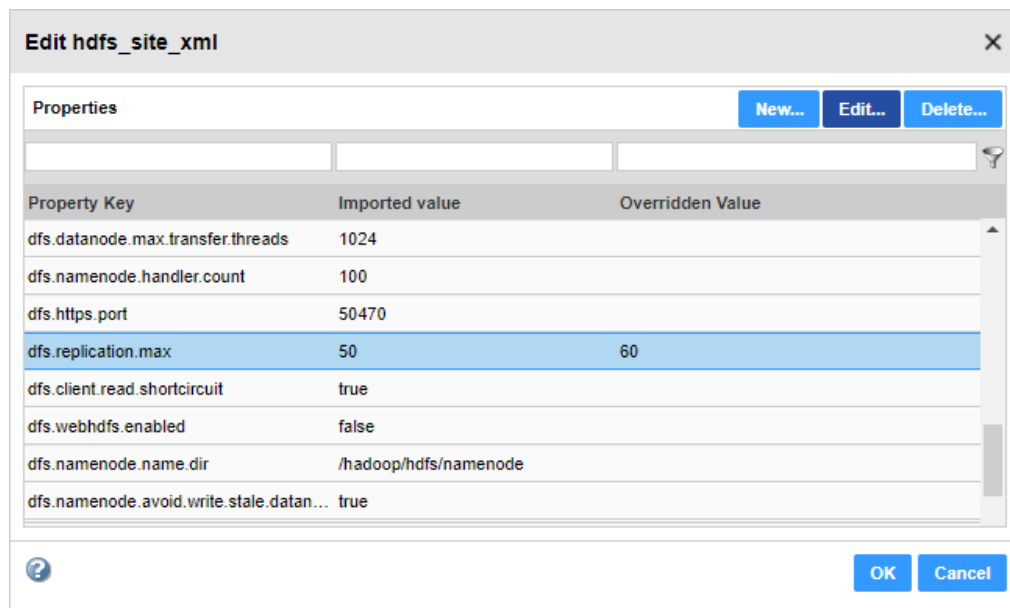
Overriding Imported Properties

You can override property values or you can update overrides from the **Active Properties** view or the **Overridden Properties** view.

1. Expand the configuration set containing the property that you want to edit.
2. Click the **Edit** icon on the name bar of the configuration set that you want to edit.
The **Edit** dialog box opens.
3. Optionally, use the filter controls to find a property.
4. Select the property to edit and click **Edit**.
The **Edit Property** dialog box opens.

5. Enter a value in the **Overridden Value** pane, and click **OK**.

The following image shows the overridden value in the **Edit** dialog box:

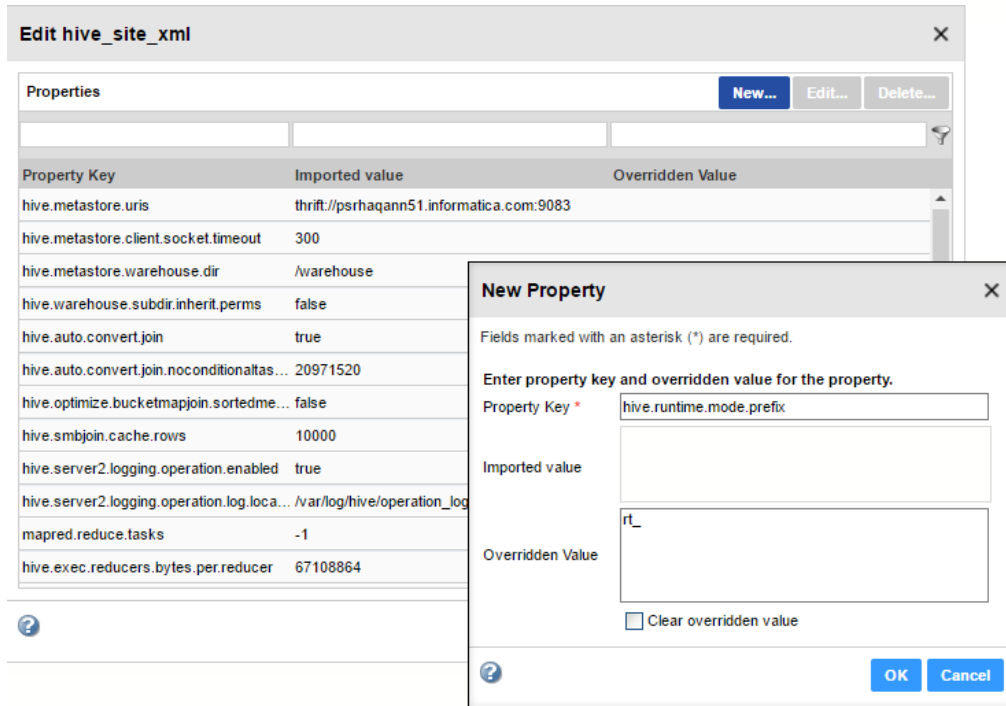


Creating User-Defined Properties

You can create user-defined properties in the **Active Properties** view or the **Overridden Properties** view. When you create a user-defined property, you configure an overridden value. You cannot configure an imported value in a user-defined property.

You can create a user-defined property based on your requirements.

1. Expand the configuration set where you want to create a property.
2. Click the **Edit** icon on the name bar of the configuration set that you want to edit. The **Edit** dialog box opens.
3. Click **New**. The **New Property** dialog box opens.



4. Configure the following properties:

Property	Description
Property Key	Name of the property that you want to enter.
Overridden Value	The property value. To clear the contents of this field, select Clear overridden value .

Important: If you create a property with the same name as a property that exists in a different configuration set, the Data Integration Service might use either property at run time, leading to unpredictable results.

5. Click **OK**.

Deleting Cluster Configuration Properties

You can delete imported and user-defined properties from a configuration set.

1. Select a cluster configuration to edit.
2. In the **Active Properties** view or the **Overridden Properties** view, expand a configuration set.
3. Click the **Edit** icon on the name bar of the configuration set that you want to edit.
The configuration set expands to show its contents.
4. Optionally, use the filter control at the top of the **Property Key** column to filter the properties.
5. Select the property and click **Delete**.

Note: Imported properties that you delete will be restored if you refresh the cluster configuration.

Refresh the Cluster Configuration

When property values change on the cluster, refresh the cluster configuration to import the changes. Similar to when you import the cluster configuration, you can refresh it either directly or from a .zip or .tar archive file.

You can refresh the cluster configuration from the Actions menu.

The refresh operation updates the cluster configuration in the following ways:

- Unedited values of imported properties are refreshed with the value from the cluster.
- The refresh operation refreshes the contents of the cluster configuration and properties in the connections that are associated with it, such as `fs.defaultFS`. For example, if you refresh a Hadoop cluster configuration from an archive file that does not contain one of the required `*-site.xml` files, the refresh cluster configuration drops the configuration set. If a missing `*-site.xml` file contained properties that get propagated to a connection, the refresh operation succeeds, but the connection refresh fails.
- If you override an imported property value, the refresh operation refreshes the imported value and maintains the overridden value. The Administrator tool displays the updated imported value and the active override value.
- If you override an imported property value, and the property is subsequently removed from the cluster, the refresh operation converts the property to a user-defined property.
- User-defined properties that do not exist in the cluster are not affected.
- If you configure a user-defined property that is subsequently configured on the cluster, the refresh operation converts the user-defined property to an imported property. The user-defined property value becomes an overridden value of the imported property.
- If you deleted an imported property, the refresh operation imports the property again if it exists on the cluster.

Note: If you refresh the cluster configuration while a mapping task is queued, the queued task uses outdated cluster configuration settings instead of the refreshed settings. This might happen when you manually refresh the cluster configuration, or start a task that automatically refreshes the cluster configuration. For best results, verify that all mapping tasks are complete before you refresh the cluster configuration.

Example - Cluster Configuration Refresh

The following example shows the process for importing a property added during a refresh operation after you create it as a user-defined property.

1. You add the following user-defined property to the cluster configuration:

```
<property>
  <name>runtime.engine.prefix</name>
  <value>rt_</value>
  <description>prefixes the runtime output table by this string</description>
</property>
```

2. The Hadoop administrator adds the property to the cluster, but with a different value, as follows:

```
<property>
  <name>runtime.engine.prefix</name>
  <value>runtime_</value>
  <description>prefixes the runtime output table by this string</description>
</property>
```

3. You refresh the cluster configuration, and the refresh operation performs the following tasks.
 - a. Converts the user-defined property to an imported property.

- b. Maintains the user-defined "rt_" value as an overridden value.
- c. Imports the cluster value "runtime_" as the imported value.

Delete a Cluster Configuration

You cannot delete a cluster configuration that has associated connections. You can associate the connections with a different cluster configuration before you delete the cluster configuration.

You can also use the `infacmd` cluster `DeleteConfiguration` command to delete connections when you delete the cluster configuration.

You must have write permission to delete a cluster configuration, and Manage Connections privilege to delete connections.

To delete a cluster configuration, click the **Actions** menu and select **Delete**.

Cluster Configuration Privileges and Permissions

You manage user security with privileges, roles, and permissions. Privileges determine the actions that users can perform on a cluster configuration. Permissions define the level of access that users and groups have to a cluster configuration.

Privileges and Roles

Cluster configuration privileges and roles determine the actions that users can perform using the Administrator tool and the `infacmd` command line program.

The following privileges and roles are required to perform certain actions on the cluster configuration:

Domain Administration privilege group

A user assigned the Administrator role for the domain can configure cluster configurations.

Manage Connections privilege

Users or groups assigned the Manage Connections privilege can create, refresh, and delete cluster configurations. Users can also set and clear configuration properties.

Permissions

Permissions control the level of access that a user or group has for a cluster configuration.

You can configure permissions for a cluster configuration in the Administrator tool and using `infacmd`.

Any cluster configuration permission that is assigned to a user or group in one tool also applies in the other tool. For example, you grant GroupA permission on ConfigurationA using the Informatica command line interface. GroupA has permission on ConfigurationA in the Developer tool also.

The following Informatica components use the cluster configuration permissions:

- Administrator tool. Enforces read, write, execute, and grant permissions on cluster configurations.

- Informatica command line interface. Enforces read, write, execute, and grant permissions on cluster configurations.
- Developer tool. Enforces read, write, and execute permissions on cluster configurations.
- Data Integration Service. Enforces execute permissions when a user tries to preview data or run a mapping, scorecard, or profile.

Types of Cluster Configuration Permissions

You can assign different permission types to users to perform the following actions:

Permission Type	Action
Read	View the cluster configuration.
Write	Edit and refresh the cluster configuration. Set and clear configuration properties. Export the cluster configuration with sensitive properties. Delete the cluster configuration. Users with write permission inherit read permission.
Execute	Run mappings in the Hadoop environment.
Grant	Grant permission on the cluster configuration to other users and groups. Users with grant permission inherit read permission.
All	Inherit read, write, execute, and grant permissions.
None	Remove permissions for the user.

Note the following default permissions for cluster configurations:

- The domain administrator has all permissions on all cluster configurations.
- The user that creates a cluster configuration has read, write, execute, and grant permission for the cluster configuration.
- All users have permission to view the cluster configuration name.

CHAPTER 6

Cloud Provisioning Configuration

This chapter includes the following topics:

- [Cloud Provisioning Configuration Overview, 59](#)
- [Verify Prerequisites , 60](#)
- [AWS Cloud Provisioning Configuration Properties, 61](#)
- [Azure Cloud Provisioning Configuration Properties, 63](#)
- [Databricks Cloud Provisioning Configuration Properties, 66](#)
- [Create the Cloud Provisioning Configuration, 66](#)
- [Create a Cluster Connection, 67](#)

Cloud Provisioning Configuration Overview

A cloud provisioning configuration is an object in the domain that contains information about the cloud platform. The cloud provisioning configuration gives the Data Integration Service the information it needs to create a cluster on the cloud platform.

Create a cloud provisioning configuration when you configure a cluster workflow. You create a cluster workflow to automate the creation of clusters and workflow tasks on an Amazon Web Services, Microsoft Azure, or the Databricks compute cluster.

The Data Integration Service uses the information in the cloud provisioning configuration to establish a relationship between the workflow Create Cluster task and the cloud platform, and to run tasks on the cluster that the workflow creates. Using authentication credentials from the cloud provisioning configuration, the Data Integration Service submits jobs to the compute cluster using the REST API.

The cluster connection that the cluster workflow uses contains a reference to the cloud provisioning configuration.

Consider the following high-level process for using the cloud provisioning connection:

1. Verify prerequisites.
2. Create the cloud provisioning configuration.
3. Create a cluster connection for the workflow.

After you create the cloud provisioning configuration and the cluster connection, a developer uses the Developer tool to create, deploy, and run a cluster workflow.

Verify Prerequisites

Verify the following cloud platform and domain prerequisites.

Cloud Platform Prerequisites

Verify the following prerequisites for the cloud platform:

- Create a user account with administrator permissions on the cloud platform.
- Create a resource on the cloud platform where you can create clusters.
 - On AWS, create a Virtual Private Cloud (VPC)
 - On Azure, create a Virtual Network (vnet)

The Informatica domain and the cluster the workflow creates must be located on this resource.

- If the Informatica domain is installed on-premises, enable DNS resolution.

Domain Prerequisites

Verify the following prerequisites for the Informatica domain:

- An Informatica domain must be installed. The domain can reside on an instance on the Amazon or Microsoft Azure cloud platform, or on an on-premises machine. If the Informatica instance is installed on-premises, you must configure the VPN to connect to the AWS VPC or Azure vnet where the cluster runs.
- You must have permissions to create connections on the domain.
- To create clusters on AWS, the AWS administrator must open the required ports for the security group to use in the VPC where you want to create the cluster.

Enable DNS Resolution from an On-Premises Informatica Domain

If the Informatica domain is installed on-premises, you must enable DNS resolution over the VPN that connects the domain and the cloud platform.

You can enable DNS resolution on Amazon AWS or Microsoft Azure.

Amazon AWS

To run mappings from an on-premises deployment of Big Data Management to a cluster on AWS, you must install and configure Unbound on an EC2 instance. Unbound enables DNS resolution from an on-premises network. To read how to install and configure Unbound in the AWS VPC, see the [AWS documentation](#).

Microsoft Azure

To run mappings from an on-premises deployment of Big Data Management to a cluster on Azure, you must use the Bind utility on the Azure virtual network.

Follow the steps in the Microsoft Azure article "[DNS Configuration](#)."

The article gives an example of the contents of the `/etc/bind/named.conf.options` file. You can put a list of available IP addresses on the domain network in the `goodclients` portion of the file. The following excerpt shows an example:

```
//Add the IP range of the joined network to this list
acl goodclients {
    1.2.3.0/24; # IP address range of the virtual network
    1.2.4.0/24;
    1.2.5.0/24;
    1.2.6.0/24;
    1.2.3.253;
    1.2.3.254;
    localhost;
```

```
    localnets;  
};
```

AWS Cloud Provisioning Configuration Properties

The properties in the AWS cloud provisioning configuration enable the Data Integration Service to contact and create resources on the AWS cloud platform.

General Properties

The following table describes cloud provisioning configuration general properties:

Property	Description
Name	Name of the cloud provisioning configuration.
ID	ID of the cloud provisioning configuration. Default: Same as the cloud provisioning configuration name.
Description.	Optional. Description of the cloud provisioning configuration.
AWS Access Key ID	Optional. ID of the AWS access key, which AWS uses to control REST or HTTP query protocol requests to AWS service APIs. If you do not specify a value, Informatica attempts to follow the Default Credential Provider Chain.
AWS Secret Access Key	Secret component of the AWS access key. Required if you specify the AWS Access Key ID.
Region	Region in which to create the cluster. This must be the region in which the VPC is running. Use AWS region values. For a list of acceptable values, see AWS documentation. Note: The region where you want to create the cluster can be different from the region in which the Informatica domain is installed.

Permissions

The following table describes cloud provisioning configuration permissions properties:

Property	Description
EMR Role	Name of the service role for the EMR cluster that you create. The role must have sufficient permissions to create a cluster, access S3 resources, and run jobs on the cluster. When the AWS administrator creates this role, they select the "EMR" role. This contains the default AmazonElasticMapReduceRole policy. You can edit the services in this policy.
EC2 Instance Profile	Name of the EC2 instance profile role that controls permissions on processes that run on the cluster. When the AWS administrator creates this role, they select the "EMR Role for EC2" role. This includes S3 access by default.
Auto Scaling Role	Required if you configure auto-scaling for the EMR cluster. This role is created when the AWS administrator configures auto-scaling on any cluster in the VPC. Default: When you leave this field blank, it is equivalent to setting the Auto Scaling role to "Proceed without role" when the AWS administrator creates a cluster in the AWS console.

EC2 Configuration

The following table describes cloud provisioning configuration EC2 configuration properties:

Property	Description
EC2 Key Pair	EC2 key pair to enable communication with the EMR cluster master node. Optional. This credential enables you to log into the cluster. Configure this property if you intend the cluster to be non-ephemeral.
EC2 Subnet	ID of the subnet on the VPC in which to create the cluster. Use the subnet ID of the EC2 instance where the cluster runs.
Master Security Group	Optional. ID of the security group for the cluster master node. Acts as a virtual firewall to control inbound and outbound traffic to cluster nodes. Security groups are created when the AWS administrator creates and configures a cluster in a VPC. In the AWS console, the property is equivalent to ElasticMapReduce-master. You can use existing security groups, or the AWS administrator might create dedicated security groups for the ephemeral cluster. If you do not specify a value, the cluster applies the default security group for the VPC.
Additional Master Security Groups	Optional. IDs of additional security groups to attach to the cluster master node. Use a comma-separated list of security group IDs.
Core and Task Security Group	Optional. ID of the security group for the cluster core and task nodes. When the AWS administrator creates and configures a cluster in the AWS console, the property is equivalent to the ElasticMapReduce-slave security group. If you do not specify a value, the cluster applies the default security group for the VPC.

Property	Description
Additional Core and Task Security Groups	Optional. IDs of additional security groups to attach to cluster core and task nodes. Use a comma-separated list of security group IDs.
Service Access Security Group	EMR managed security group for service access. Required when you provision an EMR cluster in a private subnet.

Azure Cloud Provisioning Configuration Properties

The properties in the Azure cloud provisioning configuration enable the Data Integration Service to contact and create resources on the Azure cloud platform.

Authentication Details

The following table describes authentication properties to configure:

Property	Description
Name	Name of the cloud provisioning configuration.
ID	ID of the cloud provisioning configuration. Default: Same as the cloud provisioning configuration name.
Description	Optional. Description of the cloud provisioning configuration.
Subscription ID	ID of the Azure account to use in the cluster creation process.
Tenant ID	A GUID string associated with the Azure Active Directory.
Client ID	A GUID string that is the same as the Application ID associated with the Service Principal. The Service Principal must be assigned to a role that has permission to create resources in the subscription that you identified in the Subscription ID property.
Client Secret	An octet string that provides a key associated with the client ID.

Storage Account Details

Choose to configure access to one of the following storage types:

- Azure Data Lake Storage (ADLS). See [Azure documentation](#).
- An Azure Storage Account, known as general or blob storage. See [Azure documentation](#).

The following table describes the information you need to configure Azure Data Lake Storage (ADLS) with the HDInsight cluster:

Property	Description
Azure Data Lake Store Name	Name of the ADLS storage to access. The ADLS storage and the cluster to create must reside in the same region.
Data Lake Service Principal Client ID	A credential that enables programmatic access to ADLS storage. Enables the Informatica domain to communicate with ADLS and run commands and mappings on the HDInsight cluster. The service principal is an Azure user that meets the following requirements: <ul style="list-style-type: none"> - Permissions to access required directories in ADLS storage. - Certificate-based authentication for ADLS storage. - Key-based authentication for ADLS storage.
Data Lake Service Principal Certificate Contents	The Base64 encoded text of the public certificate used with the service principal. Leave this property blank when you create the cloud provisioning configuration. After you save the cloud provisioning configuration, log in to the VM where the Informatica domain is installed and run <code>infacmd ccps updateADLSCertificate</code> to populate this property.
Data Lake Service Principal Certificate Password	Private key for the service principal. This private key must be associated with the service principal certificate.
Data Lake Service Principal Client Secret	An octet string that provides a key associated with the service principal.
Data Lake Service Principal OAUTH Token Endpoint	Endpoint for OAUTH token based authentication.

The following table describes the information you need to configure Azure General Storage, also known as blob storage, with the HDInsight cluster:

Property	Description
Azure Storage Account Name	Name of the storage account to access. Get the value from the Storage Accounts node in the Azure web console. The storage and the cluster to create must reside in the same region.
Azure Storage Account Key	A key to authenticate access to the storage account. To get the value from the Azure web console, select the storage account, then Access Keys. The console displays the account keys.

Cluster Deployment Details

The following table describes the cluster deployment properties that you configure:

Property	Description
Resource Group	Resource group in which to create the cluster. A resource group is a logical set of Azure resources.
Virtual Network Resource Group	Optional. Resource group to which the virtual network belongs. If you do not specify a resource group, the Data Integration Service assumes that the virtual network is a member of the same resource group as the cluster.
Virtual Network	Name of the virtual network or vnet where you want to create the cluster. Specify a vnet that resides in the resource group that you specified in the Virtual Network Resource Group property. The vnet must be in the same region as the region in which to create the cluster.
Subnet Name	Subnet in which to create the cluster. The subnet must be a part of the vnet that you designated in the previous property. Each vnet can have one or more subnets. The Azure administrator can choose an existing subnet or create one for the cluster.

External Hive Metastore Details

You can specify the properties to enable the cluster to connect to a Hive metastore database that is external to the cluster.

You can use an external relational database like MySQL or Amazon RDS as the Hive metastore database. The external database must be on the same cloud platform as the cluster to create.

If you do not specify an existing external database in this dialog box, the cluster creates its own database on the cluster. This database is terminated when the cluster is terminated.

The following table describes the Hive metastore database properties that you configure:

Property	Description
Database Name	Name of the Hive metastore database.
Database Server Name	Server on which the database resides. Note: The database server name on the Azure web console commonly includes the suffix <code>database.windows.net</code> . For example: <code>server123xyz.database.windows.net</code> . You can specify the database server name without the suffix and Informatica will automatically append the suffix. For example, you can specify <code>server123xyz</code> .
Database User Name	User name of the account for the domain to use to access the database.
Database Password	Password for the user account.

Databricks Cloud Provisioning Configuration Properties

The properties in the Databricks cloud provisioning configuration enable the Data Integration Service to contact and create resources on the Databricks cloud platform.

The following table describes the Databricks cloud provisioning configuration properties:

Property	Description
Name	Name of the cloud provisioning configuration.
ID	The cluster ID of the Databricks cluster.
Description	Optional description of the cloud provisioning configuration.
Databricks domain	Domain name of the Databricks deployment.
Databricks token ID	The token ID created within Databricks required for authentication. Note: If the token has an expiration date, verify that you get a new token from the Databricks administrator before it expires.

Create the Cloud Provisioning Configuration

Create the cloud provisioning configuration and configure it with information that the domain needs to access and create resources on the cloud platform.

The properties to configure in the cloud provisioning configuration depend on the cloud platform.

1. From the **Connections** tab, right-click the Domain node and select **New > Connection**.
The **New Connection** dialog box opens.
2. Choose one of the following cloud provisioning configuration types:
 - AWS Cloud Provisioning Configuration. For AWS cloud provisioning properties and values, see [“AWS Cloud Provisioning Configuration Properties” on page 61](#).
 - Azure Cloud Provisioning Configuration. For Azure cloud provisioning properties and values, see [“Azure Cloud Provisioning Configuration Properties” on page 63](#).
 - Databricks Cloud Provisioning Configuration. For Databricks cloud provisioning properties and values, see [“Databricks Cloud Provisioning Configuration Properties” on page 66](#).
3. Enter property values in the configuration wizard, then click **Finish** to create the cloud provisioning configuration.

The cloud provisioning configuration appears in the list of connections in the **Domain Navigator**.

Complete the Azure Cloud Provisioning Configuration

When you want to access Azure Data Lake Storage (ADLS) with the cluster workflow, complete the following steps after you configure and save the cloud provisioning configuration for Azure:

1. Log in to the VM where the Informatica domain is installed, and open a command shell.

2. From the command line, issue the following command:

```
/infacmd.sh ccps updateADLSCertificate -dn <domain name> -un <user name> -pd  
<password> -cpid <cloud provisioning connection name>  
-certPath <domain location of certificate>
```

The command automatically populates the Data Lake Service Principal Certificate Contents property of the cloud provisioning connection.

Create a Cluster Connection

Create a dedicated cluster connection, such as Hadoop or Databricks, to use with the cluster workflow.

The cluster connection saves property values for the Data Integration Service to use for cluster workflow operations. When you run a cluster workflow, the Data Integration Service uses settings in the cluster connection to run jobs on the cluster.

When you configure a Hadoop connection for the cluster workflow, populate the Cloud Provisioning Configuration property with the name of the cloud provisioning configuration you created for the workflow. Leave the Cluster Configuration property blank.

When you create the workflow, populate the Connection Name property of the Create Cluster task with this cluster connection.

CHAPTER 7

Data Integration Service Processing

This chapter includes the following topics:

- [Overview of Data Integration Service Processing, 68](#)
- [Data Integration Service Queueing, 69](#)
- [Big Data Job Recovery, 72](#)
- [Tuning for Big Data Processing, 74](#)

Overview of Data Integration Service Processing

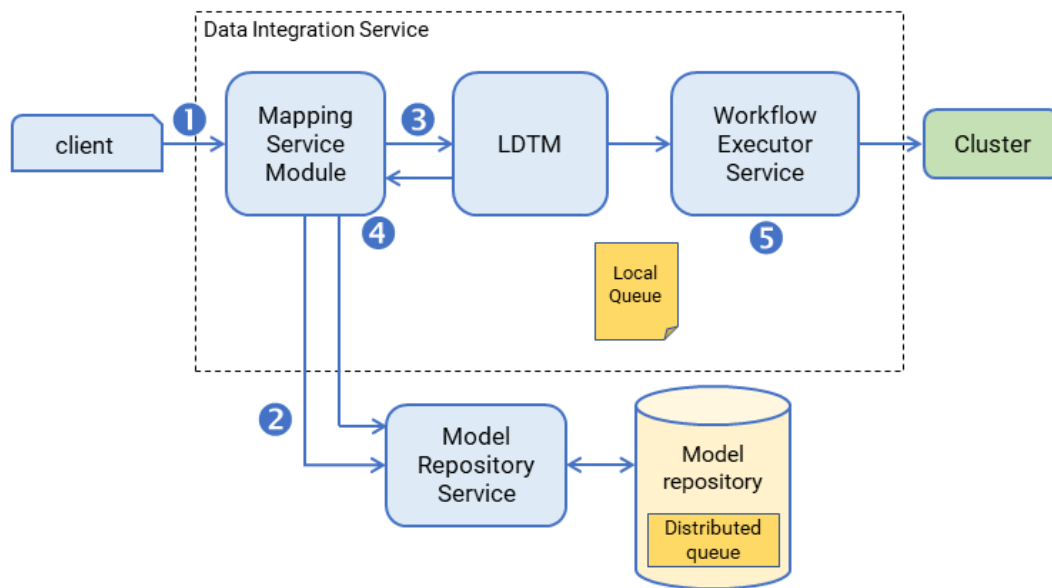
The Data Integration Service runs mappings, workflows, and other jobs that a developer deploys from clients. The Data Integration Service uses an internal process to assess the job, create an execution plan, and process the job as a series of tasks. Depending on how you configure the Data Integration Service, the process runs the tasks internally in native mode, or it sends the job for processing to a compute cluster.

When you run a mapping or other job on the Data Integration Service, the Data Integration Service saves the mapping request in a queue. The Data Integration Service mapping service takes the job from the queue when nodes are available to run the job, and creates an execution workflow. The Data Integration Service then processes the execution workflow natively or sends workflow tasks to the cluster for processing.

When a compute cluster receives a job from the Data Integration Service, it assigns each job and each child task a YARN ID. The Data Integration Service stores the YARN ID in the Model repository database to aid in tracking jobs.

The job is complete when the client receives the Complete status from the Data Integration Service.

The following image shows an overview of Data Integration Service processing:



1. A client submits a mapping execution request to the Data Integration Service. The Mapping Service Module receives the request and stores the job in the queue.
2. The Mapping Service Module connects to the Model Repository Service to fetch mapping metadata from the Model repository.
3. The Mapping Service Module passes the mapping to the Logical Data Transformation Manager (LDTM).
4. The LDTM compiles the mapping and generates the Spark execution workflow. It stores the execution workflow in the Model repository.
5. The LDTM pushes the execution workflow through the Workflow Executor Service to the cluster for processing.

You can tune the Data Integration Service and run-time engines for big data processing to ensure that sufficient resources are available to perform jobs.

Data Integration Service Queueing

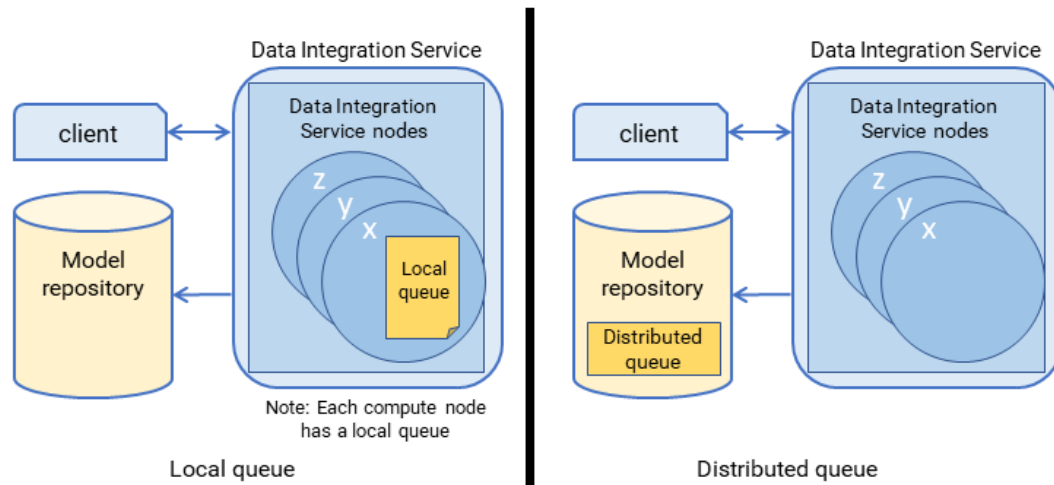
When you run a mapping or other job, the Data Integration Service stores information about the job in a queue before submitting the job for processing. Queueing protects against data loss when the node that is processing a job fails.

When you run a mapping job or workflow mapping task, the Data Integration Service adds the job to the queue. The job state appears as "Queued" in the Administrator tool contents panel. When resources are available, the Data Integration Service takes a job from the queue and runs it.

The queueing process uses the following types of queues:

- Local queue. A file-based queue on an individual compute node. Each Data Integration Service compute node contains a local queue.
- Distributed queue. A database-based queue stored in the Model repository database and accessible from any available Data Integration Service node. If the Data Integration Service is enabled for big data recovery, it uses this queue.

The following image shows both types of queue and their location:



When you run a job that cannot be queued, the Data Integration Service immediately starts running the job. If there are not enough resources available, the job fails, and you must run the job again when resources are available.

The following jobs cannot be queued:

- Jobs that cannot be deployed, such as previews and profiles
- On-demand jobs
- SQL queries
- Web service requests

Execution Pools

You can configure execution pool sizes to limit the number of jobs that can run concurrently. You can configure native, Hadoop, and on-demand execution pool sizes.

The execution pool size acts as a threshold. When the number of job requests exceeds the threshold, the Data Integration Service stops accepting new jobs.

When the number of concurrent jobs exceeds the native and Hadoop execution pool threshold, the Data Integration Service stops taking jobs from the queue until the number of concurrently running jobs falls below the threshold you configured.

When the number of concurrent jobs exceeds the on-demand execution pool threshold, the Data Integration Service rejects on-demand job requests. When the number of concurrently running jobs falls below the threshold you configured, you can resume running on-demand jobs and re-run rejected jobs.

Configure execution pool sizes in Data Integration Service Execution Options settings.

Local Queues

The local queue is stored as a file in the file system of the individual Data Integration System node that receives the request. The queue is unavailable to backup nodes or to other nodes in the grid.

The Data Integration Service uses a local queue when database-based queuing is disabled.

Note: You must use a custom command option to disable the default database queuing and enable local queuing. Contact Informatica Global Customer Support for more information.

When database-based queuing is disabled, the Data Integration Service uses the following queuing process for local queues:

1. A client submits a job request to the Data Integration Service, which stores job metadata in the local queue.
2. When the Data Integration Service node has available resources, the Data Integration Service retrieves job information from the queue and sends the next job for processing based on the order it was received in the queue.
3. If a node fails while running a job and the Data Integration Service runs on a grid or has backup nodes, the job runs on another available node, but the queue remains on the node and is unavailable to other nodes. The node must be restarted before the Data Integration can add or take new jobs from the queue.
4. When the node is restarted, the Data Integration Service picks the next job from the local queue and runs it.

When resources become available, the Data Integration Service resumes running jobs in the queue and you can run additional jobs.

Note: If the node fails, all running jobs are marked with a status of "Unknown," and you must manually re-run the jobs.

In some cases, queue files become corrupted. This might happen if disk space becomes full or the Data Integration Service node shuts down unexpectedly while writing to the file. If files are corrupted, the Data Integration Service attempts to recover them. If the files cannot be recovered, delete them. The Data Integration Service automatically creates new queue files in the node.

You might want to delete the queue files after changing Data Integration Service options, or if the Data Integration Service has not run in some time and the queue is stale.

Distributed Queues

The Data Integration Service uses a distributed queue to store job information when big data recovery is enabled. The distributed queue is stored in the Model repository and is shared by all other nodes in the grid.

Any available node can run a job from the queue. If there is a long-running job on a node, another node can pick up jobs that are waiting to run.

Consider the following queuing process for distributed queues:

1. A client submits a job request to the Data Integration Service, which stores job metadata in the distributed queue.
2. When the Data Integration Service node has available resources, the Data Integration Service retrieves the job from the queue and sends it to the available node for processing.
3. If a node fails while running a job, the job can fail over to another node. Any back-up node or node in the grid can take jobs from the queue.
4. The interrupted job runs on the new node.

You can use the command `infacmd ms abortAllJobs` to abort all jobs in the queue, or `infacmd ms purgeDatabaseWorkTables` to clear the queue.

Big Data Job Recovery

The Data Integration Service manages jobs that are deployed to run in a cluster environment. When you enable the Data Integration Service for big data recovery, the Data Integration Service can recover Hadoop mapping jobs that run on the Spark engine.

To use big data recovery, you must configure jobs to run on the Spark engine. You must configure the Data Integration Service and log settings. You must run the job from `infacmd`.

The Data Integration Service maintains a queue of jobs to run. The Data Integration Service assigns jobs from the queue to nodes which prepare them and send them to a compute cluster for processing.

The cluster assigns a YARN ID to each job and each of its child tasks to track jobs as it runs them. The Data Integration Service gets the YARN IDs from the cluster and stores them on the Model repository database.

If the Data Integration Service runs on a grid or multiple nodes, when a node fails, the Service Manager fails over to another node. The Data Integration Service queries the cluster for the status of tasks as identified by their YARN IDs and compares the response with the status of failed over tasks. Depending on the status, the Data Integration Service takes the following actions:

- If a task has no YARN ID, it submits the task to the cluster.
- If a task that has a YARN ID has not been sent to the cluster, it submits the task for processing.
- If all tasks have been sent, it continues to monitor communications from the cluster until completion.

If the Data Integration Service runs on a single node, it attempts job recovery when the node is restored.

When the Data Integration Service restarts and runs a job, the job creates a cluster configuration under the `disTemp` directory. This process causes the `disTemp` directory to grow over time. Manage disk space by monitoring and periodically clearing the contents of the `disTemp` directory.

Note: The Data Integration Service begins the recovery process by verifying that inactive nodes are not available, and then it assigns the recovered job to an available node. The verification process for unavailable nodes might take several minutes before the job is reassigned to an available node.

Scenarios Where Recovery is Possible

The Data Integration Service can recover queued or running jobs upon node failure in the following scenarios:

The job has been taken from the queue but is not yet submitted to the cluster.

Depending on the timing of node failure, the status might be Queued or Running. The Data Integration Service fails the job over to the restored node or another node and continues processing it from the point of interruption.

The job is partially submitted to the cluster.

The status of tasks sent to the cluster is Running, and the status of tasks in the queue remains Queued. The Data Integration Service identifies unsent tasks and sends them to the restored node or another node, which processes the unsent tasks and sends them to the cluster for processing.

The job is fully submitted to the cluster.

The Data Integration Service has completed its job processing, and node failure at this point has no effect on the queue or on Hadoop processing. The Data Integration Service fails over to a restored node or another node to continue communication with Hadoop and to monitor job status.

Note: If the Data Integration Service runs on a single node, the recovery process begins when the node is running again.

Scenarios Where Recovery is Not Possible

Job recovery is not possible in the following scenarios:

- The node shuts down before the Mapping Service Module stores the job request in the Model repository. The job status depends on the timing of the failure. Either the job does not appear at all, or the status is Unknown.
- You cancel a mapping while it is running. The job status depends on the timing of the cancellation request. The status will be Canceled or Complete.
- You recycle the Data Integration Service while a job is running. The job status will be Aborted or Canceled.
- The cluster fails, or communication with the cluster is lost while a job is running. The job status will be Failed.

In each case, the job must be manually re-run.

When a job has a status of Unknown, check the Data Integration Service log and check for the existence of a session log for more information.

Note: Depending on the size of the loan on the Data Integration Service grid, logs might not be updated with failed job states for several minutes.

Recovery Job Management

When you manually interrupt the Data Integration Service, the big data recovery jobs might not be aborted, and they might not be removed from the Model repository queue. Manual interruption includes disabling, recycling, or deleting the service. You can manage jobs associated with manual interruption through `infacmd`.

Consider the following ways to manage recovery jobs when you recycle, disable, or delete the Data Integration Service:

Abort jobs in the queue

If you recycle or disable the Data Integration Service on a grid while jobs are running, the service might attempt recovery when it is enabled.

If you do not want to recover the jobs, run `infacmd ms abortAllJobs` before you recycle or disable the service.

Note: If you recycle or disable a Data Integration Service on a single node or on primary and backup nodes, the service aborts all jobs and does not attempt recovery when it is enabled.

Purge jobs in the queue

When you recycle or delete a Data Integration Service, the queue in the Model repository retains job information stored for big data recovery.

To maintain the queue and purge the jobs, run `infacmd ms PurgeDatabaseWorktables`.

For more information about the commands, see the *Informatica Command Reference*.

Monitoring Recovered Jobs

You can identify recovered jobs on the Monitoring tab.

When a job is recovered, the Monitoring tab shows the same start and end time for the job, and the elapsed time = 0. While this statistic is not the actual elapsed time, it enables you to identify jobs that were recovered.

For a more accurate view of the elapsed time for the job, view the Spark job logs on the cluster or the session logs on the Data Integration Service.

Note: When the Data Integration Service recovers a job, the Administrator tool might display incomplete job statistics in the Monitoring tab when the job is complete. For example, the job statistics might not correctly display the number of rows processed.

Big Data Job Recovery Configuration

Configure a Data Integration Service for big data job recovery. If the Data Integration Service runs on a grid, verify that certain directories are configured on a network file system.

Perform the following tasks:

Enable big data job recovery.

For the Data Integration Service to recover big data jobs, you must enable the option in the Data Integration Service Execution Options.

Verify that directories are configured on a network file system.

If the Data Integration Service runs on a grid, some directories on each node must be configured on a network file system to permit the Data Integration Service components to communicate information about mapping jobs during node failures.

Verify that the following properties are configured for directories:

- Directory for temporary files. Default is `<home directory>/disTemp`
- Log directory. Default is `<Informatica installation directory>/logs/<node name>/services/DataIntegrationService`

Note: Directories must be mounted on a network file system accessible to all Data Integration Service nodes.

Tuning for Big Data Processing

Tune the application services and run-time engines for big data processing.

You might want to tune the application services and run-time engines for big data processing to ensure that the application services and the run-time engines are allocated enough resources to perform jobs.

For example, the Model Repository Service and the Data Integration Service require resources to store run-time data. When you run mappings, you might deploy the mappings to the Data Integration Service and run the mappings on the Blaze engine. Similarly, the Blaze engine requires resources to run the mappings. You must allocate enough resources between the Model Repository Service, the Data Integration Service, and the Blaze engine to ensure that mapping performance is optimized.

You can tune the application services and run-time engines based on deployment type. A deployment type represents big data processing requirements based on concurrency and volume. The deployment type defines the amount of resources that application services and run-time engines require to function efficiently, and how resources should be allocated between the application services and run-time engines.

To tune the application services and run-time engines, assess the deployment type that best describes the environment that you use for big data processing. Then select the application services and the run-time engines that you want to tune. Tune the application services and the run-time engines using `infacmd autotune autotune`.

Deployment Types

A deployment type represents big data processing requirements based on concurrency and volume.

The deployment type defines the amount of resources that application services and run-time engines require to function efficiently, and how resources should be allocated between the application services and run-time engines. The deployment types are Sandbox, Basic, Standard, and Advanced.

The following table describes the deployment types:

Deployment Type	Description
Sandbox	Used for proof of concepts or as a sandbox with minimal users.
Basic	Used for low volume processing with low levels of concurrency.
Standard	Used for high volume processing with low levels of concurrency.
Advanced	Used for high volume processing with high levels of concurrency.

Each deployment type is described using deployment criteria. The deployment criteria is a set of characteristics that are common to each deployment type. Use the deployment criteria to help you understand the deployment type that best fits the environment that you use for big data processing.

The following table defines the deployment criteria:

Deployment Criteria	Sandbox	Basic	Standard	Advanced
Total data volume	2-10 GB	10 GB - 2 TB	2 TB - 50 TB	50 TB+
Number of nodes in the Hadoop environment	2 - 5	5 - 25	25 - 100	100+
Number of developers	2	5	10	50
Number of concurrent jobs in the Hadoop environment	< 10	< 250	< 500	2000+
Number of Model repository objects	<1000	< 5000	5000 - 20000	20000+
Number of deployed applications	< 10	< 25	< 100	< 500
Number of objects per deployed application	< 10	< 50	< 100	< 100

For example, you estimate that your environment handles an average of 400 concurrent jobs and a data volume of 35 TB. According to the deployment criteria, the deployment type that best describes your environment is Standard.

Tuning the Application Services

Tune the application services for big data processing.

You tune the application services according to the deployment type that best describes the big data processing requirements in your environment. For each application service, the heap memory is tuned based on the deployment type.

The following table describes how the heap memory is tuned for each application service based on the deployment type:

Service	Sandbox	Basic	Standard	Advanced
Analyst Service	768 MB	1 GB	2 GB	4 GB
Content Management Service	1 GB	2 GB	4 GB	4 GB
Data Integration Service	640 MB	2 GB	4 GB	6 GB
Model Repository Service	1 GB	1 GB	2 GB	4 GB
Resource Manager Service	512 MB	512 MB	2 GB	4 GB
Search Service	768 MB	1 GB	2 GB	4 GB

Data Integration Service

When you tune the Data Integration Service, the deployment type additionally defines the execution pool size for jobs that run in the native and Hadoop environments.

The following table lists the execution pool size that is tuned in the native and Hadoop environments based on the deployment type:

Run-time Environment	Sandbox	Basic	Standard	Advanced
Native	10	10	15	30
Hadoop	10	500	1000	2000

Note: If the deployment type is Advanced, the Data Integration Service is tuned to run on a grid.

Tuning the Hadoop Run-time Engines

Tune the Blaze and Spark engines based on the deployment type. You tune the Blaze and Spark engines to adhere to big data processing requirements.

Tuning the Spark Engine

Tune the Spark engine according to a deployment type that defines the big data processing requirements on the Spark engine. When you tune the Spark engine, the autotune command configures the Spark advanced properties in the Hadoop connection.

The following table describes the advanced properties that are tuned:

Property	Description
spark.driver.memory	The driver process memory that the Spark engine uses to run mapping jobs.
spark.executor.memory	The amount of memory that each executor process uses to run tasklets on the Spark engine.
spark.executor.cores	The number of cores that each executor process uses to run tasklets on the Spark engine.
spark.sql.shuffle.partitions	The number of partitions that the Spark engine uses to shuffle data to process joins or aggregations in a mapping job.

The following table lists the tuned value for each advanced property based on the deployment type:

Property	Sandbox	Basic	Standard	Advanced
spark.driver.memory	1 GB	2 GB	4 GB	4 GB
spark.executor.memory	2 GB	4 GB	6 GB	6 GB
spark.executor.cores	2	2	2	2
spark.sql.shuffle.partitions	100	400	1500	3000

Tuning the Blaze Engine

Tune the Blaze engine to adhere to big data processing requirements on the Blaze engine. When you tune the Blaze engine, the autotune command configures the Blaze advanced properties in the Hadoop connection.

The following table describes the Blaze properties that are tuned:

Property	Description	Value
infagrid.orch.scheduler.oop.container.pref.memory	The amount of memory that the Blaze engine uses to run tasklets.	5120
infagrid.orch.scheduler.oop.container.pref.vcore	The number of DTM instances that run on the Blaze engine.	4
infagrid.tasklet.dtm.buffer.block.size	The amount of buffer memory that a DTM instance uses to move a block of data in a tasklet.	6553600
* The tuned properties do not depend on the deployment type.		

Autotune

Configures services and connections with recommended settings based on the deployment type. Changes take effect after you recycle the services.

For each specified service, the changes to the service take effect on all nodes that are currently configured to run the service, and the changes affect all service processes.

The infacmd autotune Autotune command uses the following syntax:

```
Autotune
<-DomainName|-dn> domain_name
<-UserName|-un> user_name
<-Password|-pd> password
[<-SecurityDomain|-sdn> security_domain]
[<-ResilienceTimeout|-re> timeout_period_in_seconds]
<-Size|-s> tuning_size_name
[<-ServiceNames|-sn> service_names]
[<-BlazeConnectionNames|-bcn> connection_names]
[<-SparkConnectionNames|-scn> connection_names]
[<-All|-a> yes_or_no]
```

The following table describes infacmd autotune Autotune options and arguments:

Option	Argument	Description
-DomainName -dn	domain_name	Required. Name of the Informatica domain. You can set the domain name with the -dn option or the environment variable INFA_DEFAULT_DOMAIN. If you set a domain name with both methods, the -dn option takes precedence.
-UserName -un	user_name	Required if the domain uses Native or LDAP authentication. User name to connect to the domain. You can set the user name with the -un option or the environment variable INFA_DEFAULT_DOMAIN_USER. If you set a user name with both methods, the -un option takes precedence. Optional if the domain uses Kerberos authentication. To run the command with single sign-on, do not set the user name. If you set the user name, the command runs without single sign-on.
-Password -pd	password	Required if you specify the user name. Password for the user name. The password is case sensitive. You can set a password with the -pd option or the environment variable INFA_DEFAULT_DOMAIN_PASSWORD. If you set a password with both methods, the password set with the -pd option takes precedence.

Option	Argument	Description
-SecurityDomain -sdn	security_domain	Required if the domain uses LDAP authentication. Optional if the domain uses native authentication or Kerberos authentication. Name of the security domain to which the domain user belongs. You can set a security domain with the -sdn option or the environment variable INFA_DEFAULT_SECURITY_DOMAIN. If you set a security domain name with both methods, the -sdn option takes precedence. The security domain name is case sensitive. If the domain uses native or LDAP authentication, the default is Native. If the domain uses Kerberos authentication, the default is the LDAP security domain created during installation. The name of the security domain is the same as the user realm specified during installation.
-ResilienceTimeout -re	timeout_period_in_seconds	Optional. Amount of time in seconds that infacmd attempts to establish or re-establish a connection to the domain. You can set the resilience timeout period with the -re option or the environment variable INFA_CLIENT_RESILIENCE_TIMEOUT. If you set the resilience timeout period with both methods, the -re option takes precedence.
-Size -s	tuning_size_name	Required. The deployment type that represents big data processing requirements based on concurrency and volume. You can enter Sandbox, Basic, Standard, or Advanced.
-ServiceNames -sn	service_names	Optional. List of services configured in the Informatica domain. Separate each service name with a comma. You can tune the following services: <ul style="list-style-type: none"> - Analyst Service - Content Management Service - Data Integration Service - Model Repository Service - Resource Manager Service - Search Service Default is none.
-BlazeConnectionNames -bcn	connection_names	Optional. List of Hadoop connections configured in the Informatica domain. For each Hadoop connection, the command tunes Blaze configuration properties in the Hadoop connection. Separate each Hadoop connection name with a comma. Default is none.

Option	Argument	Description
-SparkConnectionNames -scn	connection_names	<p>Optional. List of Hadoop connections configured in the Informatica domain. For each Hadoop connection, the command tunes Spark configuration properties in the Hadoop connection.</p> <p>Separate each Hadoop connection name with a comma.</p> <p>Default is none.</p>
-All -a	yes_or_no	<p>Optional. Enter <code>yes</code> to apply recommended settings to all Analyst Services, Content Management Services, Data Integration Services, Model Repository Services, Resource Manager Services, Search Services, and Hadoop connections in the Informatica domain.</p> <p>Enter <code>no</code> to apply the recommended settings only to the services and Hadoop connections that you specify.</p> <p>Default is <code>no</code>.</p>

APPENDIX A

Connections

This appendix includes the following topics:

- [Connections, 82](#)
- [Cloud Provisioning Configuration, 82](#)
- [Amazon Redshift Connection Properties, 88](#)
- [Amazon S3 Connection Properties, 89](#)
- [Cassandra Connection Properties, 91](#)
- [Databricks Connection Properties, 92](#)
- [Google Analytics Connection Properties, 94](#)
- [Google BigQuery Connection Properties, 94](#)
- [Google Cloud Spanner Connection Properties, 95](#)
- [Google Cloud Storage Connection Properties, 96](#)
- [Hadoop Connection Properties, 97](#)
- [HDFS Connection Properties, 102](#)
- [HBase Connection Properties, 104](#)
- [HBase Connection Properties for MapR-DB, 104](#)
- [Hive Connection Properties, 105](#)
- [JDBC Connection Properties, 108](#)
- [Kafka Connection Properties, 113](#)
- [Microsoft Azure Blob Storage Connection Properties, 114](#)
- [Microsoft Azure Cosmos DB SQL API Connection Properties, 115](#)
- [Microsoft Azure Data Lake Store Connection Properties, 116](#)
- [Microsoft Azure SQL Data Warehouse Connection Properties, 117](#)
- [Snowflake Connection Properties, 118](#)
- [Creating a Connection to Access Sources or Targets, 119](#)
- [Creating a Hadoop Connection, 119](#)
- [Configuring Hadoop Connection Properties, 121](#)

Connections

Create a connection to access non-native environments, Hadoop and Databricks. If you access HBase, HDFS, or Hive sources or targets in the Hadoop environment, you must also create those connections. You can create the connections using the Developer tool, Administrator tool, and infacmd.

You can create the following types of connections:

Hadoop connection

Create a Hadoop connection to run mappings in the Hadoop environment.

HBase connection

Create an HBase connection to access HBase. The HBase connection is a NoSQL connection.

HDFS connection

Create an HDFS connection to read data from or write data to the HDFS file system on a Hadoop cluster.

Hive connection

Create a Hive connection to access Hive as a source or target. You can access Hive as a source if the mapping is enabled for the native or Hadoop environment. You can access Hive as a target if the mapping runs on the Blaze engine.

JDBC connection

Create a JDBC connection and configure Sqoop properties in the connection to import and export relational data through Sqoop.

Databricks connection

Create a Databricks connection to run mappings in the Databricks environment.

Note: For information about creating connections to other sources or targets such as social media web sites or Teradata, see the respective PowerExchange adapter user guide for information.

Cloud Provisioning Configuration

The cloud provisioning configuration establishes a relationship between the Create Cluster task and the cluster connection that the workflows use to run mapping tasks. The Create Cluster task must include a reference to the cloud provisioning configuration. In turn, the cloud provisioning configuration points to the cluster connection that you create for use by the cluster workflow.

The properties to populate depend on the Hadoop distribution you choose to build a cluster on. Choose one of the following connection types:

- AWS Cloud Provisioning. Connects to an Amazon EMR cluster on Amazon Web Services.
- Azure Cloud Provisioning. Connects to an HDInsight cluster on the Azure platform.
- Databricks Cloud Provisioning. Connects to a Databricks cluster on the Azure Databricks platform.

AWS Cloud Provisioning Configuration Properties

The properties in the AWS cloud provisioning configuration enable the Data Integration Service to contact and create resources on the AWS cloud platform.

General Properties

The following table describes cloud provisioning configuration general properties:

Property	Description
Name	Name of the cloud provisioning configuration.
ID	ID of the cloud provisioning configuration. Default: Same as the cloud provisioning configuration name.
Description.	Optional. Description of the cloud provisioning configuration.
AWS Access Key ID	Optional. ID of the AWS access key, which AWS uses to control REST or HTTP query protocol requests to AWS service APIs. If you do not specify a value, Informatica attempts to follow the Default Credential Provider Chain.
AWS Secret Access Key	Secret component of the AWS access key. Required if you specify the AWS Access Key ID.
Region	Region in which to create the cluster. This must be the region in which the VPC is running. Use AWS region values. For a list of acceptable values, see AWS documentation. Note: The region where you want to create the cluster can be different from the region in which the Informatica domain is installed.

Permissions

The following table describes cloud provisioning configuration permissions properties:

Property	Description
EMR Role	Name of the service role for the EMR cluster that you create. The role must have sufficient permissions to create a cluster, access S3 resources, and run jobs on the cluster. When the AWS administrator creates this role, they select the "EMR" role. This contains the default AmazonElasticMapReduceRole policy. You can edit the services in this policy.
EC2 Instance Profile	Name of the EC2 instance profile role that controls permissions on processes that run on the cluster. When the AWS administrator creates this role, they select the "EMR Role for EC2" role. This includes S3 access by default.
Auto Scaling Role	Required if you configure auto-scaling for the EMR cluster. This role is created when the AWS administrator configures auto-scaling on any cluster in the VPC. Default: When you leave this field blank, it is equivalent to setting the Auto Scaling role to "Proceed without role" when the AWS administrator creates a cluster in the AWS console.

EC2 Configuration

The following table describes cloud provisioning configuration EC2 configuration properties:

Property	Description
EC2 Key Pair	EC2 key pair to enable communication with the EMR cluster master node. Optional. This credential enables you to log into the cluster. Configure this property if you intend the cluster to be non-ephemeral.
EC2 Subnet	ID of the subnet on the VPC in which to create the cluster. Use the subnet ID of the EC2 instance where the cluster runs.
Master Security Group	Optional. ID of the security group for the cluster master node. Acts as a virtual firewall to control inbound and outbound traffic to cluster nodes. Security groups are created when the AWS administrator creates and configures a cluster in a VPC. In the AWS console, the property is equivalent to ElasticMapReduce-master. You can use existing security groups, or the AWS administrator might create dedicated security groups for the ephemeral cluster. If you do not specify a value, the cluster applies the default security group for the VPC.
Additional Master Security Groups	Optional. IDs of additional security groups to attach to the cluster master node. Use a comma-separated list of security group IDs.
Core and Task Security Group	Optional. ID of the security group for the cluster core and task nodes. When the AWS administrator creates and configures a cluster in the AWS console, the property is equivalent to the ElasticMapReduce-slave security group. If you do not specify a value, the cluster applies the default security group for the VPC.
Additional Core and Task Security Groups	Optional. IDs of additional security groups to attach to cluster core and task nodes. Use a comma-separated list of security group IDs.
Service Access Security Group	EMR managed security group for service access. Required when you provision an EMR cluster in a private subnet.

Azure Cloud Provisioning Configuration Properties

The properties in the Azure cloud provisioning configuration enable the Data Integration Service to contact and create resources on the Azure cloud platform.

Authentication Details

The following table describes authentication properties to configure:

Property	Description
Name	Name of the cloud provisioning configuration.
ID	ID of the cloud provisioning configuration. Default: Same as the cloud provisioning configuration name.
Description	Optional. Description of the cloud provisioning configuration.

Property	Description
Subscription ID	ID of the Azure account to use in the cluster creation process.
Tenant ID	A GUID string associated with the Azure Active Directory.
Client ID	A GUID string that is the same as the Application ID associated with the Service Principal. The Service Principal must be assigned to a role that has permission to create resources in the subscription that you identified in the Subscription ID property.
Client Secret	An octet string that provides a key associated with the client ID.

Storage Account Details

Choose to configure access to one of the following storage types:

- Azure Data Lake Storage (ADLS). See [Azure documentation](#).
- An Azure Storage Account, known as general or blob storage. See [Azure documentation](#).

The following table describes the information you need to configure Azure Data Lake Storage (ADLS) with the HDInsight cluster:

Property	Description
Azure Data Lake Store Name	Name of the ADLS storage to access. The ADLS storage and the cluster to create must reside in the same region.
Data Lake Service Principal Client ID	A credential that enables programmatic access to ADLS storage. Enables the Informatica domain to communicate with ADLS and run commands and mappings on the HDInsight cluster. The service principal is an Azure user that meets the following requirements: <ul style="list-style-type: none"> - Permissions to access required directories in ADLS storage. - Certificate-based authentication for ADLS storage. - Key-based authentication for ADLS storage.
Data Lake Service Principal Certificate Contents	The Base64 encoded text of the public certificate used with the service principal. Leave this property blank when you create the cloud provisioning configuration. After you save the cloud provisioning configuration, log in to the VM where the Informatica domain is installed and run <code>infacmd ccps updateADLSCertificate</code> to populate this property.
Data Lake Service Principal Certificate Password	Private key for the service principal. This private key must be associated with the service principal certificate.
Data Lake Service Principal Client Secret	An octet string that provides a key associated with the service principal.
Data Lake Service Principal OAUTH Token Endpoint	Endpoint for OAUTH token based authentication.

The following table describes the information you need to configure Azure General Storage, also known as blob storage, with the HDInsight cluster:

Property	Description
Azure Storage Account Name	Name of the storage account to access. Get the value from the Storage Accounts node in the Azure web console. The storage and the cluster to create must reside in the same region.
Azure Storage Account Key	A key to authenticate access to the storage account. To get the value from the Azure web console, select the storage account, then Access Keys. The console displays the account keys.

Cluster Deployment Details

The following table describes the cluster deployment properties that you configure:

Property	Description
Resource Group	Resource group in which to create the cluster. A resource group is a logical set of Azure resources.
Virtual Network Resource Group	Optional. Resource group to which the virtual network belongs. If you do not specify a resource group, the Data Integration Service assumes that the virtual network is a member of the same resource group as the cluster.
Virtual Network	Name of the virtual network or vnet where you want to create the cluster. Specify a vnet that resides in the resource group that you specified in the Virtual Network Resource Group property. The vnet must be in the same region as the region in which to create the cluster.
Subnet Name	Subnet in which to create the cluster. The subnet must be a part of the vnet that you designated in the previous property. Each vnet can have one or more subnets. The Azure administrator can choose an existing subnet or create one for the cluster.

External Hive Metastore Details

You can specify the properties to enable the cluster to connect to a Hive metastore database that is external to the cluster.

You can use an external relational database like MySQL or Amazon RDS as the Hive metastore database. The external database must be on the same cloud platform as the cluster to create.

If you do not specify an existing external database in this dialog box, the cluster creates its own database on the cluster. This database is terminated when the cluster is terminated.

The following table describes the Hive metastore database properties that you configure:

Property	Description
Database Name	Name of the Hive metastore database.
Database Server Name	Server on which the database resides. Note: The database server name on the Azure web console commonly includes the suffix <code>database.windows.net</code> . For example: <code>server123xyz.database.windows.net</code> . You can specify the database server name without the suffix and Informatica will automatically append the suffix. For example, you can specify <code>server123xyz</code> .
Database User Name	User name of the account for the domain to use to access the database.
Database Password	Password for the user account.

Databricks Cloud Provisioning Configuration Properties

The properties in the Databricks cloud provisioning configuration enable the Data Integration Service to contact and create resources on the Databricks cloud platform.

The following table describes the Databricks cloud provisioning configuration properties:

Property	Description
Name	Name of the cloud provisioning configuration.
ID	The cluster ID of the Databricks cluster.
Description	Optional description of the cloud provisioning configuration.
Databricks domain	Domain name of the Databricks deployment.
Databricks token ID	The token ID created within Databricks required for authentication. Note: If the token has an expiration date, verify that you get a new token from the Databricks administrator before it expires.

Amazon Redshift Connection Properties

When you set up an Amazon Redshift connection, you must configure the connection properties.

The following table describes the Amazon Redshift connection properties:

Property	Description
Name	The name of the connection. The name is not case sensitive and must be unique within the domain. You can change this property after you create the connection. The name cannot exceed 128 characters, contain spaces, or contain the following special characters:~`!\$%^&*()-+={[] \:;'"<, >. ? /
ID	String that the Data Integration Service uses to identify the connection. The ID is not case sensitive. It must be 255 characters or less and must be unique in the domain. You cannot change this property after you create the connection. Default value is the connection name.
Description	The description of the connection. The description cannot exceed 4,000 characters.
Location	The domain where you want to create the connection.
Type	The connection type. Select Amazon Redshift in the Database.

The **Details** tab contains the connection attributes of the Amazon Redshift connection. The following table describes the connection attributes:

Property	Description
Username	User name of the Amazon Redshift account.
Password	Password for the Amazon Redshift account.
Schema	Optional. Amazon Redshift schema name. Do not specify the schema name if you want to use multiple schema. The Data Object wizard displays all the user-defined schemas available for the Amazon Redshift objects. Default is public.
AWS Access Key ID	Amazon S3 bucket access key ID. Note: Required if you do not use AWS Identity and Access Management (IAM) authentication.
AWS Secret Access Key	Amazon S3 bucket secret access key ID. Note: Required if you do not use AWS Identity and Access Management (IAM) authentication.
Master Symmetric Key	Optional. Provide a 256-bit AES encryption key in the Base64 format when you enable client-side encryption. You can generate a key using a third-party tool. If you specify a value, ensure that you specify the encryption type as client side encryption in the advanced target properties.

Property	Description
Customer Master Key ID	<p>Optional. Specify the customer master key ID or alias name generated by AWS Key Management Service (AWS KMS). You must generate the customer master key corresponding to the region where Amazon S3 bucket resides. You can specify any of the following values:</p> <p>Customer generated customer master key</p> <p>Enables client-side or server-side encryption.</p> <p>Default customer master key</p> <p>Enables client-side or server-side encryption. Only the administrator user of the account can use the default customer master key ID to enable client-side encryption.</p> <p>Note: You can use customer master key ID when you run a mapping in the native environment or on the Spark engine.</p>
Cluster Node Type	<p>Node type of the Amazon Redshift cluster.</p> <p>You can select the following options:</p> <ul style="list-style-type: none"> - ds1.xlarge - ds1.8xlarge - dc1.large - dc1.8xlarge - ds2.xlarge - ds2.8xlarge <p>For more information about nodes in the cluster, see the Amazon Redshift documentation.</p>
Number of Nodes in Cluster	<p>Number of nodes in the Amazon Redshift cluster.</p> <p>For more information about nodes in the cluster, see the Amazon Redshift documentation.</p>
JDBC URL	Amazon Redshift connection URL.

Note: If you upgrade the mappings created in versions 10.1.1 Update 2 or earlier, you must select the relevant schema in the connection property. Else, the mappings fail when you run them on current version.

Amazon S3 Connection Properties

When you set up an Amazon S3 connection, you must configure the connection properties.

The following table describes the Amazon S3 connection properties:

Property	Description
Name	The name of the connection. The name is not case sensitive and must be unique within the domain. You can change this property after you create the connection. The name cannot exceed 128 characters, contain spaces, or contain the following special characters:~ ` ! \$ % ^ & * () - + = { [] \ ; " ' < , > . ? /
ID	String that the Data Integration Service uses to identify the connection. The ID is not case sensitive. It must be 255 characters or less and must be unique in the domain. You cannot change this property after you create the connection. Default value is the connection name.
Description	Optional. The description of the connection. The description cannot exceed 4,000 characters.

Property	Description
Location	The domain where you want to create the connection.
Type	The Amazon S3 connection type.
Access Key	The access key ID for access to Amazon account resources. Note: Required if you do not use AWS Identity and Access Management (IAM) authentication.
Secret Key	The secret access key for access to Amazon account resources. The secret key is associated with the access key and uniquely identifies the account. Note: Required if you do not use AWS Identity and Access Management (IAM) authentication.
Folder Path	The complete path to Amazon S3 objects. The path must include the bucket name and any folder name. Do not use a slash at the end of the folder path. For example, <bucket name>/<my folder name>.
Master Symmetric Key	Optional. Provide a 256-bit AES encryption key in the Base64 format when you enable client-side encryption. You can generate a master symmetric key using a third-party tool.
Customer Master Key ID	Optional. Specify the customer master key ID or alias name generated by AWS Key Management Service (AWS KMS). You must generate the customer master key for the same region where Amazon S3 bucket reside. You can specify any of the following values: Customer generated customer master key Enables client-side or server-side encryption. Default customer master key Enables client-side or server-side encryption. Only the administrator user of the account can use the default customer master key ID to enable client-side encryption. Note: Applicable when you run a mapping in the native environment or on the Spark engine.
Region Name	Select the AWS region in which the bucket you want to access resides. Select one of the following regions: <ul style="list-style-type: none"> - Asia Pacific (Mumbai) - Asia Pacific (Seoul) - Asia Pacific (Singapore) - Asia Pacific (Sydney) - Asia Pacific (Tokyo) - AWS GovCloud (US) - Canada (Central) - China (Beijing) - China (Ningxia) - EU (Ireland) - EU (Frankfurt) - EU (London) - EU (Paris) - South America (Sao Paulo) - US East (Ohio) - US East (N. Virginia) - US West (N. California) - US West (Oregon) Default is US East (N. Virginia).

Cassandra Connection Properties

When you set up a Cassandra connection, you must configure the connection properties.

Note: The order of the connection properties might vary depending on the tool where you view them.

The following table describes the Cassandra connection properties:

Property	Description
Name	The name of the connection. The name is not case sensitive and must be unique within the domain. You can change this property after you create the connection. The name cannot exceed 128 characters, contain spaces, or contain the following special characters: ~ ` ! \$ % ^ & * () - + = { [] \ : ; " ' < , > . ? /
ID	String that the Data Integration Service uses to identify the connection. The ID is not case sensitive. The ID must be 255 characters or less and must be unique in the domain. You cannot change this property after you create the connection. Default value is the connection name.
Description	Optional. The description of the connection. The description cannot exceed 4,000 characters.
Location	The domain where you want to create the connection.
Type	The connection type. Select Cassandra .
Host Name	Host name or IP address of the Cassandra server.
Port	Cassandra server port number. Default is 9042.
User Name	User name to access the Cassandra server.
Password	Password corresponding to the user name to access the Cassandra server.
Default Keyspace	Name of the Cassandra keyspace to use by default.
SQL Identifier Character	Type of character that the database uses to enclose delimited identifiers in SQL or CQL queries. The available characters depend on the database type. Select None if the database uses regular identifiers. When the Data Integration Service generates SQL or CQL queries, the service does not place delimited characters around any identifiers. Select a character if the database uses delimited identifiers. When the Data Integration Service generates SQL or CQL queries, the service encloses delimited identifiers within this character.
Additional Connection Properties	Enter one or more JDBC connection parameters in the following format: <param1>=<value>;<param2>=<value>;<param3>=<value> PowerExchange for Cassandra JDBC supports the following JDBC connection parameters: <ul style="list-style-type: none"> - BinaryColumnLength - DecimalColumnScale - EnableCaseSensitive - EnableNullInsert - EnablePaging - RowsPerPage - StringColumnLength - VTableSeparator

Property	Description
SSL Mode	Not applicable for PowerExchange for Cassandra JDBC. Select disabled .
SSL Truststore Path	Not applicable for PowerExchange for Cassandra JDBC.
SSL Truststore Password	Not applicable for PowerExchange for Cassandra JDBC.
SSL Keystore Path	Not applicable for PowerExchange for Cassandra JDBC.
SSL Keystore Password	Not applicable for PowerExchange for Cassandra JDBC.

Databricks Connection Properties

Use the Databricks connection to run mappings on a Databricks cluster.

A Databricks connection is a cluster type connection. You can create and manage a Databricks connection in the Administrator tool or the Developer tool. You can use infacmd to create a Databricks connection. Configure properties in the Databricks connection to enable communication between the Data Integration Service and the Databricks cluster.

The following table describes the general connection properties for the Databricks connection:

Property	Description
Name	The name of the connection. The name is not case sensitive and must be unique within the domain. You can change this property after you create the connection. The name cannot exceed 128 characters, contain spaces, or contain the following special characters:~ `!\$%^&*() - + = {[]} \ \ ; " ' < , > . ? /
ID	String that the Data Integration Service uses to identify the connection. The ID is not case sensitive. It must be 255 characters or less and must be unique in the domain. You cannot change this property after you create the connection. Default value is the connection name.
Description	Optional. The description of the connection. The description cannot exceed 4,000 characters.
Connection Type	Choose Databricks.
Cluster Configuration	Name of the cluster configuration associated with the Databricks environment. Required if you do not configure the cloud provisioning configuration.
Cloud Provisioning Configuration	Name of the cloud provisioning configuration associated with a Databricks cloud platform. Required if you do not configure the cluster configuration.

Property	Description
Staging Directory	The directory where the Databricks Spark engine stages run-time files. If you specify a directory that does not exist, the Data Integration Service creates it at run time. If you do not provide a directory path, the run-time staging files are written to <code></cluster staging directory>/DATABRICKS</code> .
Advanced Properties	List of advanced properties that are unique to the Databricks environment. You can configure run-time properties for the Databricks environment in the Data Integration Service and in the Databricks connection. You can override a property configured at a high level by setting the value at a lower level. For example, if you configure a property in the Data Integration Service custom properties, you can override it in the Databricks connection. The Data Integration Service processes property overrides based on the following priorities: 1. Databricks connection advanced properties 2. Data Integration Service custom properties Note: Informatica does not recommend changing these property values before you consult with third-party documentation, Informatica documentation, or Informatica Global Customer Support. If you change a value without knowledge of the property, you might experience performance degradation or other unexpected results.

Advanced Properties

Configure the following properties in the **Advanced Properties** of the Databricks configuration section:

infaspark.json.parser.mode

Specifies the parser how to handle corrupt JSON records. You can set the value to one of the following modes:

- **DROPMALFORMED.** The parser ignores all corrupted records. Default mode.
- **PERMISSIVE.** The parser accepts non-standard fields as nulls in corrupted records.
- **FAILFAST.** The parser generates an exception when it encounters a corrupted record and the Spark application goes down.

infaspark.json.parser.multiLine

Specifies whether the parser can read a multiline record in a JSON file. You can set the value to true or false. Default is false. Applies only to non-native distributions that use Spark version 2.2.x and above.

infaspark.flatfile.writer.nullValue

When the Databricks Spark engine writes to a target, it converts null values to empty strings (" "). For example, 12, AB,"",23p09udj.

The Databricks Spark engine can write the empty strings to string columns, but when it tries to write an empty string to a non-string column, the mapping fails with a type mismatch.

To allow the Databricks Spark engine to convert the empty strings back to null values and write to the target, configure the following advanced property in the Databricks Spark connection:

```
infaspark.flatfile.writer.nullValue=true
```

Google Analytics Connection Properties

When you set up a Google Analytics connection, you must configure the connection properties.

Note: The order of the connection properties might vary depending on the tool where you view them.

The following table describes the Google Analytics connection properties:

Property	Description
Name	The name of the connection. The name is not case sensitive and must be unique within the domain. You can change this property after you create the connection. The name cannot exceed 128 characters, contain spaces, or contain the following special characters:~ `! \$ % ^ & * () - + = { [] \ : ; " ' < , > . ? /
ID	String that the Data Integration Service uses to identify the connection. The ID is not case sensitive. The ID must be 255 characters or less and must be unique in the domain. You cannot change this property after you create the connection. Default value is the connection name.
Description	Optional. The description of the connection. The description cannot exceed 4,000 characters.
Location	The domain where you want to create the connection.
Type	The connection type. Select Google Analytics .
Service Account ID	Specifies the client_email value present in the JSON file that you download after you create a service account.
Service Account Key	Specifies the private_key value present in the JSON file that you download after you create a service account.
APIVersion	API that PowerExchange for Google Analytics uses to read from Google Analytics reports. Select Core Reporting API v3 . Note: PowerExchange for Google Analytics does not support Analytics Reporting API v4.

Google BigQuery Connection Properties

When you set up a Google BigQuery connection, you must configure the connection properties.

Note: The order of the connection properties might vary depending on the tool where you view them.

The following table describes the Google BigQuery connection properties:

Property	Description
Service Account ID	Specifies the client_email value present in the JSON file that you download after you create a service account in Google BigQuery.
Service Account Key	Specifies the private_key value present in the JSON file that you download after you create a service account in Google BigQuery.

Property	Description
Connection mode	<p>The mode that you want to use to read data from or write data to Google BigQuery.</p> <p>Select one of the following connection modes:</p> <ul style="list-style-type: none"> - Simple. Flattens each field within the Record data type field as a separate field in the mapping. - Hybrid. Displays all the top-level fields in the Google BigQuery table including Record data type fields. PowerExchange for Google BigQuery displays the top-level Record data type field as a single field of the String data type in the mapping. - Complex. Displays all the columns in the Google BigQuery table as a single field of the String data type in the mapping. <p>Default is Simple.</p>
Schema Definition File Path	<p>Specifies a directory on the client machine where the PowerCenter Integration ServiceData Integration Service must create a JSON file with the sample schema of the Google BigQuery table. The JSON file name is the same as the Google BigQuery table name.</p> <p>Alternatively, you can specify a storage path in Google Cloud Storage where the PowerCenter Integration ServiceData Integration Service must create a JSON file with the sample schema of the Google BigQuery table. You can download the JSON file from the specified storage path in Google Cloud Storage to a local machine.</p>
Project ID	<p>Specifies the project_id value present in the JSON file that you download after you create a service account in Google BigQuery.</p> <p>If you have created multiple projects with the same service account, enter the ID of the project that contains the dataset that you want to connect to.</p>
Storage Path	<p>This property applies when you read or write large volumes of data.</p> <p>Path in Google Cloud Storage where the PowerCenter Integration ServiceData Integration Service creates a local stage file to store the data temporarily.</p> <p>You can either enter the bucket name or the bucket name and folder name.</p> <p>For example, enter <code>gs://<bucket_name></code> or <code>gs://<bucket_name>/<folder_name></code></p>

Google Cloud Spanner Connection Properties

When you set up a Google Cloud Spanner connection, you must configure the connection properties.

Note: The order of the connection properties might vary depending on the tool where you view them.

The following table describes the Google Cloud Spanner connection properties:

Property	Description
Name	<p>The name of the connection. The name is not case sensitive and must be unique within the domain. You can change this property after you create the connection. The name cannot exceed 128 characters, contain spaces, or contain the following special characters: ~ ` ! \$ % ^ & * () - + = { } \ : ; " ' < , > . ? /</p>
ID	<p>String that the Data Integration Service uses to identify the connection.</p> <p>The ID is not case sensitive. The ID must be 255 characters or less and must be unique in the domain. You cannot change this property after you create the connection.</p> <p>Default value is the connection name.</p>
Description	<p>Optional. The description of the connection. The description cannot exceed 4,000 characters.</p>

Property	Description
Location	The domain where you want to create the connection.
Type	The connection type. Select Google Cloud Spanner.
Project ID	Specifies the project_id value present in the JSON file that you download after you create a service account. If you have created multiple projects with the same service account, enter the ID of the project that contains the bucket that you want to connect to.
Service Account ID	Specifies the client_email value present in the JSON file that you download after you create a service account.
Service Account Key	Specifies the private_key value present in the JSON file that you download after you create a service account.
Instance ID	Name of the instance that you created in Google Cloud Spanner.

Google Cloud Storage Connection Properties

When you set up a Google Cloud Storage connection, you must configure the connection properties.

Note: The order of the connection properties might vary depending on the tool where you view them.

The following table describes the Google Cloud Storage connection properties:

Property	Description
Name	The name of the connection. The name is not case sensitive and must be unique within the domain. You can change this property after you create the connection. The name cannot exceed 128 characters, contain spaces, or contain the following special characters:~ `! \$ % ^ & * () - + = { } \ : ; " ' < , > . ? /
ID	String that the Data Integration Service uses to identify the connection. The ID is not case sensitive. The ID must be 255 characters or less and must be unique in the domain. You cannot change this property after you create the connection. Default value is the connection name.
Description	Optional. The description of the connection. The description cannot exceed 4,000 characters.
Location	The domain where you want to create the connection.
Type	The connection type. Select Google Cloud Storage .
Project ID	Specifies the project_id value present in the JSON file that you download after you create a service account. If you have created multiple projects with the same service account, enter the ID of the project that contains the bucket that you want to connect to.

Property	Description
Service Account ID	Specifies the client_email value present in the JSON file that you download after you create a service account.
Service Account Key	Specifies the private_key value present in the JSON file that you download after you create a service account.

Hadoop Connection Properties

Use the Hadoop connection to configure mappings to run on a Hadoop cluster. A Hadoop connection is a cluster type connection. You can create and manage a Hadoop connection in the Administrator tool or the Developer tool. You can use `infacmd` to create a Hadoop connection. Hadoop connection properties are case sensitive unless otherwise noted.

Hadoop Cluster Properties

Configure properties in the Hadoop connection to enable communication between the Data Integration Service and the Hadoop cluster.

The following table describes the general connection properties for the Hadoop connection:

Property	Description
Name	The name of the connection. The name is not case sensitive and must be unique within the domain. You can change this property after you create the connection. The name cannot exceed 128 characters, contain spaces, or contain the following special characters: ~ ` ! \$ % ^ & * () - + = { [] } \ : ; " ' < , > . ? /
ID	String that the Data Integration Service uses to identify the connection. The ID is not case sensitive. It must be 255 characters or less and must be unique in the domain. You cannot change this property after you create the connection. Default value is the connection name.
Description	The description of the connection. Enter a string that you can use to identify the connection. The description cannot exceed 4,000 characters.
Cluster Configuration	The name of the cluster configuration associated with the Hadoop environment. Required if you do not configure the Cloud Provisioning Configuration.
Cloud Provisioning Configuration	Name of the cloud provisioning configuration associated with a cloud platform such as Amazon AWS or Microsoft Azure. Required if you do not configure the Cluster Configuration.

Property	Description
Cluster Environment Variables*	<p>Environment variables that the Hadoop cluster uses.</p> <p>For example, the variable ORACLE_HOME represents the directory where the Oracle database client software is installed.</p> <p>You can configure run-time properties for the Hadoop environment in the Data Integration Service, the Hadoop connection, and in the mapping. You can override a property configured at a high level by setting the value at a lower level. For example, if you configure a property in the Data Integration Service custom properties, you can override it in the Hadoop connection or in the mapping. The Data Integration Service processes property overrides based on the following priorities:</p> <ol style="list-style-type: none"> 1. Mapping custom properties set using infacmd ms runMapping with the -cp option 2. Mapping run-time properties for the Hadoop environment 3. Hadoop connection advanced properties for run-time engines 4. Hadoop connection advanced general properties, environment variables, and classpaths 5. Data Integration Service custom properties
Cluster Library Path*	<p>The path for shared libraries on the cluster.</p> <p>The \$DEFAULT_CLUSTER_LIBRARY_PATH variable contains a list of default directories.</p>
Cluster Classpath*	<p>The classpath to access the Hadoop jar files and the required libraries.</p> <p>The \$DEFAULT_CLUSTER_CLASSPATH variable contains a list of paths to the default jar files and libraries.</p> <p>You can configure run-time properties for the Hadoop environment in the Data Integration Service, the Hadoop connection, and in the mapping. You can override a property configured at a high level by setting the value at a lower level. For example, if you configure a property in the Data Integration Service custom properties, you can override it in the Hadoop connection or in the mapping. The Data Integration Service processes property overrides based on the following priorities:</p> <ol style="list-style-type: none"> 1. Mapping custom properties set using infacmd ms runMapping with the -cp option 2. Mapping run-time properties for the Hadoop environment 3. Hadoop connection advanced properties for run-time engines 4. Hadoop connection advanced general properties, environment variables, and classpaths 5. Data Integration Service custom properties
Cluster Executable Path*	<p>The path for executable files on the cluster.</p> <p>The \$DEFAULT_CLUSTER_EXEC_PATH variable contains a list of paths to the default executable files.</p>

* Informatica does not recommend changing these property values before you consult with third-party documentation, Informatica documentation, or Informatica Global Customer Support. If you change a value without knowledge of the property, you might experience performance degradation or other unexpected results.

Common Properties

The following table describes the common connection properties that you configure for the Hadoop connection:

Property	Description
Impersonation User Name	<p>Required if the Hadoop cluster uses Kerberos authentication. Hadoop impersonation user. The user name that the Data Integration Service impersonates to run mappings in the Hadoop environment.</p> <p>The Data Integration Service runs mappings based on the user that is configured. Refer the following order to determine which user the Data Integration Services uses to run mappings:</p> <ol style="list-style-type: none"> 1. Operating system profile user. The mapping runs with the operating system profile user if the profile user is configured. If there is no operating system profile user, the mapping runs with the Hadoop impersonation user. 2. Hadoop impersonation user. The mapping runs with the Hadoop impersonation user if the operating system profile user is not configured. If the Hadoop impersonation user is not configured, the Data Integration Service runs mappings with the Data Integration Service user. 3. Informatica services user. The mapping runs with the operating user that starts the Informatica daemon if the operating system profile user and the Hadoop impersonation user are not configured.
Temporary Table Compression Codec	<p>Hadoop compression library for a compression codec class name.</p> <p>Note: The Spark engine does not support compression settings for temporary tables. When you run mappings on the Spark engine, the Spark engine stores temporary tables in an uncompressed file format.</p>
Codec Class Name	<p>Codec class name that enables data compression and improves performance on temporary staging tables.</p>
Hive Staging Database Name	<p>Namespace for Hive staging tables. Use the name <code>default</code> for tables that do not have a specified database name.</p> <p>If you do not configure a namespace, the Data Integration Service uses the Hive database name in the Hive target connection to create staging tables.</p> <p>When you run a mapping in the native environment to write data to Hive, you must configure the Hive staging database name in the Hive connection. The Data Integration Service ignores the value you configure in the Hadoop connection.</p>
Advanced Properties	<p>List of advanced properties that are unique to the Hadoop environment. The properties are common to the Blaze and Spark engines. The advanced properties include a list of default properties.</p> <p>You can configure run-time properties for the Hadoop environment in the Data Integration Service, the Hadoop connection, and in the mapping. You can override a property configured at a high level by setting the value at a lower level. For example, if you configure a property in the Data Integration Service custom properties, you can override it in the Hadoop connection or in the mapping. The Data Integration Service processes property overrides based on the following priorities:</p> <ol style="list-style-type: none"> 1. Mapping custom properties set using <code>infacmd ms runMapping</code> with the <code>-cp</code> option 2. Mapping run-time properties for the Hadoop environment 3. Hadoop connection advanced properties for run-time engines 4. Hadoop connection advanced general properties, environment variables, and classpaths 5. Data Integration Service custom properties <p>Note: Informatica does not recommend changing these property values before you consult with third-party documentation, Informatica documentation, or Informatica Global Customer Support. If you change a value without knowledge of the property, you might experience performance degradation or other unexpected results.</p>

Reject Directory Properties

The following table describes the connection properties that you configure to the Hadoop Reject Directory.

Property	Description
Write Reject Files to Hadoop	If you use the Blaze engine to run mappings, select the check box to specify a location to move reject files. If checked, the Data Integration Service moves the reject files to the HDFS location listed in the property, Reject File Directory. By default, the Data Integration Service stores the reject files based on the RejectDir system parameter.
Reject File Directory	The directory for Hadoop mapping files on HDFS when you run mappings.

Hive Pushdown Configuration

Note: Effective in version 10.2.2, Informatica dropped support for the Hive engine. Do not configure the pushdown properties related to the Hive engine.

Blaze Configuration

The following table describes the connection properties that you configure for the Blaze engine:

Property	Description
Blaze Staging Directory	The HDFS file path of the directory that the Blaze engine uses to store temporary files. Verify that the directory exists. The YARN user, Blaze engine user, and mapping impersonation user must have write permission on this directory. Default is <code>/blaze/workdir</code> . If you clear this property, the staging files are written to the Hadoop staging directory <code>/tmp/blaze_<user name></code> .
Blaze User Name	The owner of the Blaze service and Blaze service logs. When the Hadoop cluster uses Kerberos authentication, the default user is the Data Integration Service SPN user. When the Hadoop cluster does not use Kerberos authentication and the Blaze user is not configured, the default user is the Data Integration Service user.
Minimum Port	The minimum value for the port number range for the Blaze engine. Default is 12300.
Maximum Port	The maximum value for the port number range for the Blaze engine. Default is 12600.
YARN Queue Name	The YARN scheduler queue name used by the Blaze engine that specifies available resources on a cluster.
Blaze Job Monitor Address	The host name and port number for the Blaze Job Monitor. Use the following format: <code><hostname>:<port></code> Where - <code><hostname></code> is the host name or IP address of the Blaze Job Monitor server. - <code><port></code> is the port on which the Blaze Job Monitor listens for remote procedure calls (RPC). For example, enter: <code>myhostname:9080</code>

Property	Description
Blaze YARN Node Label	<p>Node label that determines the node on the Hadoop cluster where the Blaze engine runs. If you do not specify a node label, the Blaze engine runs on the nodes in the default partition.</p> <p>If the Hadoop cluster supports logical operators for node labels, you can specify a list of node labels. To list the node labels, use the operators && (AND), (OR), and ! (NOT).</p>
Advanced Properties	<p>List of advanced properties that are unique to the Blaze engine. The advanced properties include a list of default properties.</p> <p>You can configure run-time properties for the Hadoop environment in the Data Integration Service, the Hadoop connection, and in the mapping. You can override a property configured at a high level by setting the value at a lower level. For example, if you configure a property in the Data Integration Service custom properties, you can override it in the Hadoop connection or in the mapping. The Data Integration Service processes property overrides based on the following priorities:</p> <ol style="list-style-type: none"> 1. Mapping custom properties set using <code>infacmd ms runMapping</code> with the <code>-cp</code> option 2. Mapping run-time properties for the Hadoop environment 3. Hadoop connection advanced properties for run-time engines 4. Hadoop connection advanced general properties, environment variables, and classpaths 5. Data Integration Service custom properties <p>Note: Informatica does not recommend changing these property values before you consult with third-party documentation, Informatica documentation, or Informatica Global Customer Support. If you change a value without knowledge of the property, you might experience performance degradation or other unexpected results.</p>

Spark Configuration

The following table describes the connection properties that you configure for the Spark engine:

Property	Description
Spark Staging Directory	<p>The HDFS file path of the directory that the Spark engine uses to store temporary files for running jobs. The YARN user, Data Integration Service user, and mapping impersonation user must have write permission on this directory.</p> <p>If you do not specify a file path, by default, the temporary files are written to the Hadoop staging directory <code>/tmp/SPARK_<user name></code>.</p> <p>When you run Sqoop jobs on the Spark engine, the Data Integration Service creates a Sqoop staging directory within the Spark staging directory to store temporary files: <code><Spark staging directory>/sqoop_staging</code></p>
Spark Event Log Directory	Optional. The HDFS file path of the directory that the Spark engine uses to log events.

Property	Description
YARN Queue Name	The YARN scheduler queue name used by the Spark engine that specifies available resources on a cluster. The name is case sensitive.
Advanced Properties	<p>List of advanced properties that are unique to the Spark engine. The advanced properties include a list of default properties.</p> <p>You can configure run-time properties for the Hadoop environment in the Data Integration Service, the Hadoop connection, and in the mapping. You can override a property configured at a high level by setting the value at a lower level. For example, if you configure a property in the Data Integration Service custom properties, you can override it in the Hadoop connection or in the mapping. The Data Integration Service processes property overrides based on the following priorities:</p> <ol style="list-style-type: none"> 1. Mapping custom properties set using <code>infacmd ms runMapping</code> with the <code>-cp</code> option 2. Mapping run-time properties for the Hadoop environment 3. Hadoop connection advanced properties for run-time engines 4. Hadoop connection advanced general properties, environment variables, and classpaths 5. Data Integration Service custom properties <p>Note: Informatica does not recommend changing these property values before you consult with third-party documentation, Informatica documentation, or Informatica Global Customer Support. If you change a value without knowledge of the property, you might experience performance degradation or other unexpected results.</p>

HDFS Connection Properties

Use a Hadoop File System (HDFS) connection to access data in the Hadoop cluster. The HDFS connection is a file system type connection. You can create and manage an HDFS connection in the Administrator tool, Analyst tool, or the Developer tool. HDFS connection properties are case sensitive unless otherwise noted.

Note: The order of the connection properties might vary depending on the tool where you view them.

The following table describes HDFS connection properties:

Property	Description
Name	Name of the connection. The name is not case sensitive and must be unique within the domain. The name cannot exceed 128 characters, contain spaces, or contain the following special characters: ~ ` ! \$ % ^ & * () - + = { [] \ : ; " ' < , > . ? /
ID	String that the Data Integration Service uses to identify the connection. The ID is not case sensitive. It must be 255 characters or less and must be unique in the domain. You cannot change this property after you create the connection. Default value is the connection name.
Description	The description of the connection. The description cannot exceed 765 characters.
Location	The domain where you want to create the connection. Not valid for the Analyst tool.
Type	The connection type. Default is Hadoop File System.

Property	Description
User Name	User name to access HDFS.
NameNode URI	<p>The URI to access the storage system.</p> <p>You can find the value for <code>fs.defaultFS</code> in the <code>core-site.xml</code> configuration set of the cluster configuration.</p> <p>Note: If you create connections when you import the cluster configuration, the NameNode URI property is populated by default, and it is updated each time you refresh the cluster configuration. If you manually set this property or override the value, the refresh operation does not update this property.</p>

Accessing Multiple Storage Types

Use the NameNode URI property in the connection parameters to connect to various storage types. The following table lists the storage type and the NameNode URI format for the storage type:

Storage	NameNode URI Format
HDFS	<p><code>hdfs://<namenode>:<port></code></p> <p>where:</p> <ul style="list-style-type: none"> - <code><namenode></code> is the host name or IP address of the NameNode. - <code><port></code> is the port that the NameNode listens for remote procedure calls (RPC). <p><code>hdfs://<nameservice></code> in case of NameNode high availability.</p>
MapR-FS	<code>maprfs:///</code>
WASB in HDInsight	<p><code>wasb://<container_name>@<account_name>.blob.core.windows.net/<path></code></p> <p>where:</p> <ul style="list-style-type: none"> - <code><container_name></code> identifies a specific Azure Storage Blob container. <p>Note: <code><container_name></code> is optional.</p> <ul style="list-style-type: none"> - <code><account_name></code> identifies the Azure Storage Blob object. <p>Example:</p> <p><code>wasb://infabdmoffering1storage.blob.core.windows.net/infabdmoffering1cluster/mr-history</code></p>
ADLS in HDInsight	<code>adl://home</code>

When you create a cluster configuration from an Azure HDInsight cluster, the cluster configuration uses either ADLS or WASB as the primary storage. You cannot create a cluster configuration with ADLS or WASB as the secondary storage. You can edit the NameNode URI property in the HDFS connection to connect to a local HDFS location.

HBase Connection Properties

Use an HBase connection to access HBase. The HBase connection is a NoSQL connection. You can create and manage an HBase connection in the Administrator tool or the Developer tool. HBase connection properties are case sensitive unless otherwise noted.

The following table describes HBase connection properties:

Property	Description
Name	The name of the connection. The name is not case sensitive and must be unique within the domain. You can change this property after you create the connection. The name cannot exceed 128 characters, contain spaces, or contain the following special characters: ~ ` ! \$ % ^ & * () - + = { [] } \ : ; " ' < , > . ? /
ID	String that the Data Integration Service uses to identify the connection. The ID is not case sensitive. It must be 255 characters or less and must be unique in the domain. You cannot change this property after you create the connection. Default value is the connection name.
Description	The description of the connection. The description cannot exceed 4,000 characters.
Location	The domain where you want to create the connection.
Type	The connection type. Select HBase.
Database Type	Type of database that you want to connect to. Select HBase to create a connection for an HBase table.

HBase Connection Properties for MapR-DB

Use an HBase connection to connect to a MapR-DB table. The HBase connection is a NoSQL connection. You can create and manage an HBase connection in the Administrator tool or the Developer tool. HBase connection properties are case sensitive unless otherwise noted.

The following table describes the HBase connection properties for MapR-DB:

Property	Description
Name	Name of the connection. The name is not case sensitive and must be unique within the domain. You can change this property after you create the connection. The name cannot exceed 128 characters, contain spaces, or contain the following special characters: ~ ` ! \$ % ^ & * () - + = { [] } \ : ; " ' < , > . ? /
ID	String that the Data Integration Service uses to identify the connection. The ID is not case sensitive. It must be 255 characters or less and must be unique in the domain. You cannot change this property after you create the connection. Default value is the connection name.

Property	Description
Description	Description of the connection. The description cannot exceed 4,000 characters.
Location	Domain where you want to create the connection.
Type	Connection type. Select HBase .
Database Type	Type of database that you want to connect to. Select MapR-DB to create a connection for a MapR-DB table.
Cluster Configuration	The name of the cluster configuration associated with the Hadoop environment.
MapR-DB Database Path	Database path that contains the MapR-DB table that you want to connect to. Enter a valid MapR cluster path. When you create an HBase data object for MapR-DB, you can browse only tables that exist in the MapR-DB path that you specify in the Database Path field. You cannot access tables that are available in sub-directories in the specified path. For example, if you specify the path as <code>/user/customers/</code> , you can access the tables in the <code>customers</code> directory. However, if the <code>customers</code> directory contains a sub-directory named <code>regions</code> , you cannot access the tables in the following directory: <code>/user/customers/regions</code>

Hive Connection Properties

Use the Hive connection to access Hive data. A Hive connection is a database type connection. You can create and manage a Hive connection in the Administrator tool, Analyst tool, or the Developer tool. Hive connection properties are case sensitive unless otherwise noted.

Note: The order of the connection properties might vary depending on the tool where you view them.

The following table describes Hive connection properties:

Property	Description
Name	The name of the connection. The name is not case sensitive and must be unique within the domain. You can change this property after you create the connection. The name cannot exceed 128 characters, contain spaces, or contain the following special characters: <code>~ ` ! \$ % ^ & * () - + = { [] } \ : ; " ' < , > . ? /</code>
ID	String that the Data Integration Service uses to identify the connection. The ID is not case sensitive. It must be 255 characters or less and must be unique in the domain. You cannot change this property after you create the connection. Default value is the connection name.
Description	The description of the connection. The description cannot exceed 4000 characters.

Property	Description
Location	The domain where you want to create the connection. Not valid for the Analyst tool.
Type	The connection type. Select Hive.
LDAP username	<p>LDAP user name of the user that the Data Integration Service impersonates to run mappings on a Hadoop cluster. The user name depends on the JDBC connection string that you specify in the Metadata Connection String or Data Access Connection String for the native environment.</p> <p>If the Hadoop cluster uses Kerberos authentication, the principal name for the JDBC connection string and the user name must be the same. Otherwise, the user name depends on the behavior of the JDBC driver. With Hive JDBC driver, you can specify a user name in many ways and the user name can become a part of the JDBC URL.</p> <p>If the Hadoop cluster does not use Kerberos authentication, the user name depends on the behavior of the JDBC driver.</p> <p>If you do not specify a user name, the Hadoop cluster authenticates jobs based on the following criteria:</p> <ul style="list-style-type: none"> - The Hadoop cluster does not use Kerberos authentication. It authenticates jobs based on the operating system profile user name of the machine that runs the Data Integration Service. - The Hadoop cluster uses Kerberos authentication. It authenticates jobs based on the SPN of the Data Integration Service. LDAP username will be ignored.
Password	Password for the LDAP username.
Environment SQL	<p>SQL commands to set the Hadoop environment. In native environment type, the Data Integration Service executes the environment SQL each time it creates a connection to a Hive metastore. If you use the Hive connection to run profiles on a Hadoop cluster, the Data Integration Service executes the environment SQL at the beginning of each Hive session.</p> <p>The following rules and guidelines apply to the usage of environment SQL in both connection modes:</p> <ul style="list-style-type: none"> - Use the environment SQL to specify Hive queries. - Use the environment SQL to set the classpath for Hive user-defined functions and then use environment SQL or PreSQL to specify the Hive user-defined functions. You cannot use PreSQL in the data object properties to specify the classpath. If you use Hive user-defined functions, you must copy the .jar files to the following directory: <pre><Informatica installation directory>/services/shared/hadoop/ <Hadoop distribution name>/extras/hive-auxjars</pre> <ul style="list-style-type: none"> - You can use environment SQL to define Hadoop or Hive parameters that you want to use in the PreSQL commands or in custom queries. - If you use multiple values for the Environment SQL property, ensure that there is no space between the values.
SQL Identifier Character	The type of character used to identify special characters and reserved SQL keywords, such as WHERE. The Data Integration Service places the selected character around special characters and reserved SQL keywords. The Data Integration Service also uses this character for the Support mixed-case identifiers property.

Properties to Access Hive as Source or Target

The following table describes the connection properties that you configure to access Hive as a source or target:

Property	Description
JDBC Driver Class Name	Name of the Hive JDBC driver class. If you leave this option blank, the Developer tool uses the default Apache Hive JDBC driver shipped with the distribution. If the default Apache Hive JDBC driver does not fit your requirements, you can override the Apache Hive JDBC driver with a third-party Hive JDBC driver by specifying the driver class name.
Metadata Connection String	<p>The JDBC connection URI used to access the metadata from the Hadoop server.</p> <p>You can use PowerExchange for Hive to communicate with a HiveServer service or HiveServer2 service. To connect to HiveServer, specify the connection string in the following format:</p> <pre>jdbc:hive2://<hostname>:<port>/<db></pre> <p>Where</p> <ul style="list-style-type: none"> - <hostname> is name or IP address of the machine on which HiveServer2 runs. - <port> is the port number on which HiveServer2 listens. - <db> is the database name to which you want to connect. If you do not provide the database name, the Data Integration Service uses the default database details. <p>To connect to HiveServer2, use the connection string format that Apache Hive implements for that specific Hadoop Distribution. For more information about Apache Hive connection string formats, see the Apache Hive documentation.</p> <p>For user impersonation, you must add <code>hive.server2.proxy.user=<xyz></code> to the JDBC connection URI. If you do not configure user impersonation, the current user's credentials are used connect to the HiveServer2.</p> <p>If the Hadoop cluster uses SSL or TLS authentication, you must add <code>ssl=true</code> to the JDBC connection URI. For example: <code>jdbc:hive2://<hostname>:<port>/<db>;ssl=true</code></p> <p>If you use self-signed certificate for SSL or TLS authentication, ensure that the certificate file is available on the client machine and the Data Integration Service machine. For more information, see the <i>Informatica Big Data Management Integration Guide</i>.</p>
Bypass Hive JDBC Server	<p>JDBC driver mode. Select the check box to use the embedded JDBC driver mode.</p> <p>To use the JDBC embedded mode, perform the following tasks:</p> <ul style="list-style-type: none"> - Verify that Hive client and Informatica services are installed on the same machine. - Configure the Hive connection properties to run mappings on a Hadoop cluster. <p>If you choose the non-embedded mode, you must configure the Data Access Connection String. Informatica recommends that you use the JDBC embedded mode.</p>
Fine Grained Authorization	<p>When you select the option to observe fine grained authorization in a Hive source, the mapping observes the following:</p> <ul style="list-style-type: none"> - Row and column level restrictions. Applies to Hadoop clusters where Sentry or Ranger security modes are enabled. - Data masking rules. Applies to masking rules set on columns containing sensitive data by Dynamic Data Masking. <p>If you do not select the option, the Blaze and Spark engines ignore the restrictions and masking rules, and results include restricted or sensitive data.</p>

Property	Description
Data Access Connection String	<p>The connection string to access data from the Hadoop data store. To connect to HiveServer, specify the non-embedded JDBC mode connection string in the following format:</p> <pre>jdbc:hive2://<hostname>:<port>/<db></pre> <p>Where</p> <ul style="list-style-type: none"> - <hostname> is name or IP address of the machine on which HiveServer2 runs. - <port> is the port number on which HiveServer2 listens. - <db> is the database to which you want to connect. If you do not provide the database name, the Data Integration Service uses the default database details. <p>To connect to HiveServer2, use the connection string format that Apache Hive implements for the specific Hadoop Distribution. For more information about Apache Hive connection string formats, see the Apache Hive documentation.</p> <p>For user impersonation, you must add <code>hive.server2.proxy.user=<xyz></code> to the JDBC connection URI. If you do not configure user impersonation, the current user's credentials are used connect to the HiveServer2.</p> <p>If the Hadoop cluster uses SSL or TLS authentication, you must add <code>ssl=true</code> to the JDBC connection URI. For example: <code>jdbc:hive2://<hostname>:<port>/<db>;ssl=true</code></p> <p>If you use self-signed certificate for SSL or TLS authentication, ensure that the certificate file is available on the client machine and the Data Integration Service machine. For more information, see the <i>Informatica Big Data Management Integration Guide</i>.</p>
Hive Staging Directory on HDFS	<p>HDFS directory for Hive staging tables. You must grant execute permission to the Hadoop impersonation user and the mapping impersonation users.</p> <p>This option is applicable and required when you write data to a Hive target in the native environment.</p>
Hive Staging Database Name	<p>Namespace for Hive staging tables. Use the name <code>default</code> for tables that do not have a specified database name.</p> <p>This option is applicable when you run a mapping in the native environment to write data to a Hive target.</p> <p>If you run the mapping on the Blaze or Spark engine, you do not need to configure the Hive staging database name in the Hive connection. The Data Integration Service uses the value that you configure in the Hadoop connection.</p>

JDBC Connection Properties

You can use a JDBC connection to access tables in a database. You can create and manage a JDBC connection in the Administrator tool, the Developer tool, or the Analyst tool.

Note: The order of the connection properties might vary depending on the tool where you view them.

The following table describes JDBC connection properties:

Property	Description
Database Type	The database type.
Name	Name of the connection. The name is not case sensitive and must be unique within the domain. The name cannot exceed 128 characters, contain spaces, or contain the following special characters: ~ ` ! \$ % ^ & * () - + = { [] \ : ; " ' < , > . ? /
ID	String that the Data Integration Service uses to identify the connection. The ID is not case sensitive. It must be 255 characters or less and must be unique in the domain. You cannot change this property after you create the connection. Default value is the connection name.
Description	The description of the connection. The description cannot exceed 765 characters.
User Name	The database user name.
Password	The password for the database user name.
JDBC Driver Class Name	<p>Name of the JDBC driver class.</p> <p>The following list provides the driver class name that you can enter for the applicable database type:</p> <ul style="list-style-type: none"> - DataDirect JDBC driver class name for Oracle: <code>com.informatica.jdbc.oracle.OracleDriver</code> - DataDirect JDBC driver class name for IBM DB2: <code>com.informatica.jdbc.db2.DB2Driver</code> - DataDirect JDBC driver class name for Microsoft SQL Server: <code>com.informatica.jdbc.sqlserver.SQLServerDriver</code> - DataDirect JDBC driver class name for Sybase ASE: <code>com.informatica.jdbc.sybase.SybaseDriver</code> - DataDirect JDBC driver class name for Informix: <code>com.informatica.jdbc.informix.InformixDriver</code> - DataDirect JDBC driver class name for MySQL: <code>com.informatica.jdbc.mysql.MySQLDriver</code> <p>For more information about which driver class to use with specific databases, see the vendor documentation.</p>

Property	Description
Connection String	<p>Connection string to connect to the database. Use the following connection string:</p> <pre>jdbc:<subprotocol>:<subname></pre> <p>The following list provides sample connection strings that you can enter for the applicable database type:</p> <ul style="list-style-type: none"> - Connection string for DataDirect Oracle JDBC driver: <pre>jdbc:informatica:oracle://<host>:<port>;SID=<value></pre> - Connection string for Oracle JDBC driver: <pre>jdbc:oracle:thin:@//<host>:<port>:<SID></pre> - Connection string for DataDirect IBM DB2 JDBC driver: <pre>jdbc:informatica:db2://<host>:<port>;DatabaseName=<value></pre> - Connection string for IBM DB2 JDBC driver: <pre>jdbc:db2://<host>:<port>/<database_name></pre> - Connection string for DataDirect Microsoft SQL Server JDBC driver: <pre>jdbc:informatica:sqlserver://<host>;DatabaseName=<value></pre> - Connection string for Microsoft SQL Server JDBC driver: <pre>jdbc:sqlserver://<host>;DatabaseName=<value></pre> - Connection string for Netezza JDBC driver: <pre>jdbc:netezza://<host>:<port>/<database_name></pre> - Connection string for Pivotal Greenplum driver: <pre>jdbc:pivotal:greenplum://<host>:<port>;/database_name=<value></pre> - Connection string for Postgres Greenplum driver: <pre>jdbc:postgres://<host>:<port>/<database_name></pre> - Connection string for Teradata JDBC driver: <pre>jdbc:teradata://<host>/database_name=<value>,tmode=<value>,charset=<value></pre> <p>For more information about the connection string to use with specific drivers, see the vendor documentation.</p>
Environment SQL	<p>Optional. Enter SQL commands to set the database environment when you connect to the database. The Data Integration Service executes the connection environment SQL each time it connects to the database.</p> <p>Note: If you enable Sqoop, Sqoop ignores this property.</p>
Transaction SQL	<p>Optional. Enter SQL commands to set the database environment when you connect to the database. The Data Integration Service executes the transaction environment SQL at the beginning of each transaction.</p> <p>Note: If you enable Sqoop, Sqoop ignores this property.</p>
SQL Identifier Character	<p>Type of character that the database uses to enclose delimited identifiers in SQL queries. The available characters depend on the database type.</p> <p>Select (None) if the database uses regular identifiers. When the Data Integration Service generates SQL queries, the service does not place delimited characters around any identifiers.</p> <p>Select a character if the database uses delimited identifiers. When the Data Integration Service generates SQL queries, the service encloses delimited identifiers within this character.</p> <p>Note: If you enable Sqoop, Sqoop ignores this property.</p>
Support Mixed-case Identifiers	<p>Enable if the database uses case-sensitive identifiers. When enabled, the Data Integration Service encloses all identifiers within the character selected for the SQL Identifier Character property.</p> <p>When the SQL Identifier Character property is set to none, the Support Mixed-case Identifiers property is disabled.</p> <p>Note: If you enable Sqoop, Sqoop honors this property when you generate and execute a DDL script to create or replace a target at run time. In all other scenarios, Sqoop ignores this property.</p>

Sqoop Connection-Level Arguments

In the JDBC connection, you can define the arguments that Sqoop must use to connect to the database. The Data Integration Service merges the arguments that you specify with the default command that it constructs based on the JDBC connection properties. The arguments that you specify take precedence over the JDBC connection properties.

If you want to use the same driver to import metadata and run the mapping, and do not want to specify any additional Sqoop arguments, select **Sqoop v1.x** from the **Use Sqoop Version** list and leave the **Sqoop Arguments** field empty in the JDBC connection. The Data Integration Service constructs the Sqoop command based on the JDBC connection properties that you specify.

However, if you want to use a different driver for run-time tasks or specify additional run-time Sqoop arguments, select **Sqoop v1.x** from the **Use Sqoop Version** list and specify the arguments in the **Sqoop Arguments** field.

You can configure the following Sqoop arguments in the JDBC connection:

driver

Defines the JDBC driver class that Sqoop must use to connect to the database.

Use the following syntax:

```
--driver <JDBC driver class>
```

For example, use the following syntax depending on the database type that you want to connect to:

- **Aurora:** `--driver com.mysql.jdbc.Driver`
- **Greenplum:** `--driver org.postgresql.Driver`
- **IBM DB2:** `--driver com.ibm.db2.jcc.DB2Driver`
- **IBM DB2 z/OS:** `--driver com.ibm.db2.jcc.DB2Driver`
- **Microsoft SQL Server:** `--driver com.microsoft.sqlserver.jdbc.SQLServerDriver`
- **Netezza:** `--driver org.netezza.Driver`
- **Oracle:** `--driver oracle.jdbc.driver.OracleDriver`
- **Teradata:** `--driver com.teradata.jdbc.TeraDriver`

connect

Defines the JDBC connection string that Sqoop must use to connect to the database. The JDBC connection string must be based on the driver that you define in the driver argument.

Use the following syntax:

```
--connect <JDBC connection string>
```

For example, use the following syntax depending on the database type that you want to connect to:

- **Aurora:** `--connect "jdbc:mysql://<host_name>:<port>/<schema_name>"`
- **Greenplum:** `--connect jdbc:postgresql://<host_name>:<port>/<database_name>`
- **IBM DB2:** `--connect jdbc:db2://<host_name>:<port>/<database_name>`
- **IBM DB2 z/OS:** `--connect jdbc:db2://<host_name>:<port>/<database_name>`
- **Microsoft SQL Server:** `--connect jdbc:sqlserver://<host_name>:<port or named_instance>;databaseName=<database_name>`
- **Netezza:** `--connect "jdbc:netezza://<database_server_name>:<port>/<database_name>;schema=<schema_name>"`

- Oracle: `--connect jdbc:oracle:thin:@<database_host_name>:<database_port>:<database_SID>`
- Teradata: `--connect jdbc:teradata://<host_name>/database=<database_name>`

connection-manager

Defines the connection manager class name that Sqoop must use to connect to the database.

Use the following syntax:

```
--connection-manager <connection manager class name>
```

For example, use the following syntax to use the generic JDBC manager class name:

```
--connection-manager org.apache.sqoop.manager.GenericJdbcManager
```

direct

When you read data from or write data to Oracle, you can configure the `direct` argument to enable Sqoop to use OraOop. OraOop is a specialized Sqoop plug-in for Oracle that uses native protocols to connect to the Oracle database. When you configure OraOop, the performance improves.

You can configure OraOop when you run Sqoop mappings on the Spark engine.

Use the following syntax:

```
--direct
```

When you use OraOop, you must use the following syntax to specify multiple arguments:

```
-D<argument=value> -D<argument=value>
```

Note: If you specify multiple arguments and include a space character between `-D` and the argument name-value pair, Sqoop considers only the first argument and ignores the remaining arguments.

If you do not direct the job to a specific queue, the Spark engine uses the default queue.

-Dsqoop.connection.factories

To run the mapping on the Blaze engine with the Teradata Connector for Hadoop (TDCH) specialized connectors for Sqoop, you must configure the `-Dsqoop.connection.factories` argument. Use the argument to define the TDCH connection factory class that Sqoop must use. The connection factory class varies based on the TDCH Sqoop Connector that you want to use.

- To use Cloudera Connector Powered by Teradata, configure the `-Dsqoop.connection.factories` argument as follows:

```
-Dsqoop.connection.factories=com.cloudera.connector.teradata.TeradataManagerFactory
```
- To use Hortonworks Connector for Teradata (powered by the Teradata Connector for Hadoop), configure the `-Dsqoop.connection.factories` argument as follows:

```
-Dsqoop.connection.factories=org.apache.sqoop.teradata.TeradataManagerFactory
```

Note: To run the mapping on the Spark engine, you do not need to configure the `-Dsqoop.connection.factories` argument. The Data Integration Service invokes Cloudera Connector Powered by Teradata and Hortonworks Connector for Teradata (powered by the Teradata Connector for Hadoop) by default.

--infaoptimize

Use this argument to disable the performance optimization of Sqoop pass-through mappings on the Spark engine.

When you run a Sqoop pass-through mapping on the Spark engine, the Data Integration Service optimizes mapping performance in the following scenarios:

- You read data from a Sqoop source and write data to a Hive target that uses the Text format.
- You read data from a Sqoop source and write data to an HDFS target that uses the Flat, Avro, or Parquet format.

If you want to disable the performance optimization, set the `--infaoptimize` argument to `false`. For example, if you see data type issues after you run an optimized Sqoop mapping, you can disable the performance optimization.

Use the following syntax:

```
--infaoptimize false
```

For a complete list of the Sqoop arguments that you can configure, see the Sqoop documentation.

Kafka Connection Properties

The Kafka connection is a messaging connection. Use the Kafka connection to access Kafka as a message target. You can create and manage a Kafka connection in the Developer tool or through `infacmd`.

The following table describes the general connection properties for the Kafka connection:

Property	Description
Name	The name of the connection. The name is not case sensitive and must be unique within the domain. You can change this property after you create the connection. The name cannot exceed 128 characters, contain spaces, or contain the following special characters: ~ ` ! \$ % ^ & * () - + = { [] \ : ; " ' < , > . ? /
ID	The string that the Data Integration Service uses to identify the connection. The ID is not case sensitive. It must be 255 characters or less and must be unique in the domain. You cannot change this property after you create the connection. Default value is the connection name.
Description	The description of the connection. Enter a string that you can use to identify the connection. The description cannot exceed 4,000 characters.
Location	The domain where you want to create the connection.
Type	The connection type.

The following table describes the Kafka broker properties for the Kafka connection:

Property	Description
Kafka Broker List	Comma-separated list of Kafka brokers which maintains the configuration of the Kafka messaging broker. To specify a Kafka broker, use the following format: <IP Address>:<port>
ZooKeeper Host Port List	Optional. Comma-separated list of Apache ZooKeeper which maintains the configuration of the Kafka messaging broker. To specify the ZooKeeper, use the following format: <IP Address>:<port>
Retry Timeout	Number of seconds the Integration Service attempts to reconnect to the Kafka broker to write data. If the source or target is not available for the time you specify, the mapping execution stops to avoid any data loss.
Kafka Broker Version	Configure the Kafka messaging broker version to 0.10.1.x-2.0.0.

Microsoft Azure Blob Storage Connection Properties

Use a Microsoft Azure SQL Blob Storage connection to access a Microsoft Azure Blob Storage.

Note: The order of the connection properties might vary depending on the tool where you view them.

You can create and manage a Microsoft Azure Blob Storage connection in the Administrator tool or the Developer tool. The following table describes the Microsoft Azure Blob Storage connection properties:

Property	Description
Name	Name of the Microsoft Azure Blob Storage connection.
ID	String that the Data Integration Service uses to identify the connection. The ID is not case sensitive. It must be 255 characters or less and must be unique in the domain. You cannot change this property after you create the connection. Default value is the connection name.
Description	Description of the connection.
Location	The domain where you want to create the connection.
Type	Type of connection. Select AzureBlob.

The **Connection Details** tab contains the connection attributes of the Microsoft Azure Blob Storage connection. The following table describes the connection attributes:

Property	Description
Account Name	Name of the Microsoft Azure Storage account.
Account Key	Microsoft Azure Storage access key.
Container Name	The root container or sub-folders with the absolute path.
Endpoint Suffix	Type of Microsoft Azure end-points. You can select any of the following end-points: <ul style="list-style-type: none"> - <code>core.windows.net</code>: Default - <code>core.usgovcloudapi.net</code>: To select the US government Microsoft Azure end-points - <code>core.chinacloudapi.cn</code>: Not applicable

Microsoft Azure Cosmos DB SQL API Connection Properties

Use a Microsoft Azure Cosmos DB connection to connect to the Cosmos DB database. When you create a Microsoft Azure Cosmos DB connection, you enter information for metadata and data access.

The following table describes the Microsoft Azure Cosmos DB connection properties:

Property	Description
Name	Name of the Cosmos DB connection.
ID	String that the Data Integration Service uses to identify the connection. The ID is not case sensitive. It must be 255 characters or less and must be unique in the domain. You cannot change this property after you create the connection. Default value is the connection name.
Description	Description of the connection. The description cannot exceed 765 characters.
Location	The project or folder in the Model repository where you want to store the Cosmos DB connection.
Type	Select Microsoft Azure Cosmos DB SQL API.
Cosmos DB URI	The URI of Microsoft Azure Cosmos DB account.
Key	The primary and secondary key to which provides you complete administrative access to the resources within Microsoft Azure Cosmos DB account.
Database	Name of the database that contains the collections from which you want to read or write JSON documents.

Note: You can find the Cosmos DB URI and Key values in the **Keys** settings on Azure portal. Contact your Azure administrator for more details.

Microsoft Azure Data Lake Store Connection Properties

Use a Microsoft Azure Data Lake Store connection to access a Microsoft Azure Data Lake Store.

Note: The order of the connection properties might vary depending on the tool where you view them.

You can create and manage a Microsoft Azure SQL Data Warehouse connection in the Administrator tool or the Developer tool. The following table describes the Microsoft Azure Data Lake Store connection properties:

Property	Description
Name	The name of the connection. The name is not case sensitive and must be unique within the domain. You can change this property after you create the connection. The name cannot exceed 128 characters, contain spaces, or contain the following special characters: ~ ` ! \$ % ^ & * () - + = { [] \ : ; " ' < , > . ? /
ID	String that the Data Integration Service uses to identify the connection. The ID is not case sensitive. It must be 255 characters or less and must be unique in the domain. You cannot change this property after you create the connection. Default value is the connection name.
Description	The description of the connection. The description cannot exceed 4,000 characters.
Location	The domain where you want to create the connection.
Type	The connection type. Select Microsoft Azure Data Lake Store.

The following table describes the properties for metadata access:

Property	Description
ADLS Account Name	The name of the Microsoft Azure Data Lake Store.
ClientID	The ID of your application to complete the OAuth Authentication in the Active Directory.
Client Secret	The client secret key to complete the OAuth Authentication in the Active Directory.
Directory	The Microsoft Azure Data Lake Store directory that you use to read data or write data. The default is root directory.
AuthEndpoint	The OAuth 2.0 token endpoint from where access code is generated based on based on the Client ID and Client secret is completed.

For more information about creating a client ID, client secret, and auth end point, contact the Azure administrator or see Microsoft Azure Data Lake Store documentation.

Microsoft Azure SQL Data Warehouse Connection Properties

Use a Microsoft Azure SQL Data Warehouse connection to access a Microsoft Azure SQL Data Warehouse.

Note: The order of the connection properties might vary depending on the tool where you view them.

You can create and manage a Microsoft Azure SQL Data Warehouse connection in the Administrator tool or the Developer tool. The following table describes the Microsoft Azure SQL Data Warehouse connection properties:

Property	Description
Name	The name of the connection. The name is not case sensitive and must be unique within the domain. You can change this property after you create the connection. The name cannot exceed 128 characters, contain spaces, or contain the following special characters: ~ ` ! \$ % ^ & * () - + = { [] \ ; : " ' < , > . ? /
ID	String that the Data Integration Service uses to identify the connection. The ID is not case sensitive. It must be 255 characters or less and must be unique in the domain. You cannot change this property after you create the connection. Default value is the connection name.
Description	The description of the connection. The description cannot exceed 4,000 characters.
Location	The domain where you want to create the connection.
Type	The connection type. Select Microsoft Azure SQL Data Warehouse.

The following table describes the properties for metadata access:

Property	Description
Azure DW JDBC URL	Microsoft Azure Data Warehouse JDBC connection string. For example, you can enter the following connection string: jdbc:sqlserver:// <Server>.database.windows.net:1433;database=<Database>. The Administrator can download the URL from Microsoft Azure portal.
Azure DW JDBC Username	User name to connect to the Microsoft Azure SQL Data Warehouse account. You must have permission to read, write, and truncate data in Microsoft Azure SQL Data Warehouse.
Azure DW JDBC Password	Password to connect to the Microsoft Azure SQL Data Warehouse account.
Azure DW Schema Name	Name of the schema in Microsoft Azure SQL Data Warehouse.
Azure Blob Account Name	Name of the Microsoft Azure Storage account to stage the files.

Property	Description
Azure Blob Account Key	The key that authenticates the access to the Blob storage account.
Blob End-point	Type of Microsoft Azure end-points. You can select any of the following end-points: <ul style="list-style-type: none"> - core.windows.net: Default - core.usgovcloudapi.net: To select the US government Microsoft Azure end-points - core.chinacloudapi.cn: Not applicable You can configure the US government Microsoft Azure end-points when a mapping runs in the native environment and on the Spark engine.

Snowflake Connection Properties

When you set up a Snowflake connection, you must configure the connection properties.

Note: The order of the connection properties might vary depending on the tool where you view them.

The following table describes the Snowflake connection properties:

Property	Description
Name	The name of the connection. The name is not case sensitive and must be unique within the domain. You can change this property after you create the connection. The name cannot exceed 128 characters, contain spaces, or contain the following special characters: ~ ` ! \$ % ^ & * () - + = { [] \ ; " ' < , > . ? /
ID	String that the Data Integration Service uses to identify the connection. The ID is not case sensitive. The ID must be 255 characters or less and must be unique in the domain. You cannot change this property after you create the connection. Default value is the connection name.
Description	Optional. The description of the connection. The description cannot exceed 4,000 characters.
Location	The domain where you want to create the connection.
Type	The connection type. Select Snowflake.
Username	The user name to connect to the Snowflake account.
Password	The password to connect to the Snowflake account.
Account	The name of the Snowflake account.
Warehouse	The Snowflake warehouse name.

Property	Description
Role	The Snowflake role assigned to the user.
Additional JDBC URL Parameters	<p>Enter one or more JDBC connection parameters in the following format:</p> <pre><param1>=<value>&<param2>=<value>&<param3>=<value>...</pre> <p>For example:</p> <pre>user=jon&warehouse=mywh&db=mydb&schema=public</pre> <p>To access Snowflake through Okta SSO authentication, enter the web-based IdP implementing SAML 2.0 protocol in the following format:</p> <pre>authenticator=https://<Your_Okta_Account_Name>.okta.com</pre> <p>Note: Microsoft ADFS is not supported.</p> <p>For more information about configuring Okta authentication, see the following website: https://docs.snowflake.net/manuals/user-guide/admin-security-fed-auth-configure-snowflake.html#configuring-snowflake-to-use-federated-authentication</p>

Creating a Connection to Access Sources or Targets

Create connections before you import data objects, preview data, and profile data.

1. Within the Administrator tool click **Manage > Connections**.
2. Select **Actions > New > Connection**.
3. Select the type of connection that you want to create:
 - To select an HBase connection, select **NoSQL > HBase**.
 - To select an HDFS connection, select **File Systems > Hadoop File System**.
 - To select a Hive connection, select **Database > Hive**.
 - To select a JDBC connection, select **Database > JDBC**.
4. Click **OK**.
5. Enter a connection name, ID, and optional description.
6. Configure the connection properties. For a Hive connection, you must choose the **Access Hive as a source or target** option to use Hive as a source or a target.
7. Click **Test Connection** to verify the connection.
8. Click **Finish**.

Creating a Hadoop Connection

Create a Hadoop connection before you run a mapping in the Hadoop environment.

1. Click **Window > Preferences**.
2. Select **Informatica > Connections**.

3. Expand the domain in the **Available Connections** list.
4. Select the **Cluster** connection type in the **Available Connections** list and click **Add**.
The **New Cluster Connection** dialog box appears.
5. Enter the general properties for the connection.

New Cluster Connection

Cluster Connection
Provide the connection details.

Name:

ID:

Description:

Location:

Type:

?

6. Click **Next**.
7. Enter the Hadoop cluster properties, common properties, and the reject directory properties.
8. Click **Next**.
9. Click **Next**.
Effective in version 10.2.2, Informatica dropped support for the Hive engine. Do not enter Hive configuration properties.
10. Enter configuration properties for the Blaze engine and click **Next**.
11. Enter configuration properties for the Spark engine and click **Finish**.

Configuring Hadoop Connection Properties

When you create a Hadoop connection, default values are assigned to cluster environment variables, cluster path properties, and advanced properties. You can add or edit values for these properties. You can also reset to default values.

You can configure the following Hadoop connection properties based on the cluster environment and functionality that you use:

- Cluster Environment Variables
- Cluster Library Path
- Common Advanced Properties
- Blaze Engine Advanced Properties
- Spark Engine Advanced Properties

Note: Informatica does not recommend changing these property values before you consult with third-party documentation, Informatica documentation, or Informatica Global Customer Support. If you change a value without knowledge of the property, you might experience performance degradation or other unexpected results.

To reset to default values, delete the property values. For example, if you delete the values of an edited Cluster Library Path property, the value resets to the default \$DEFAULT_CLUSTER_LIBRARY_PATH.

Cluster Environment Variables

Cluster Environment Variables property lists the environment variables that the cluster uses. Each environment variable contains a name and a value. You can add environment variables or edit environment variables.

To edit the property in the text box, use the following format with &: to separate each name-value pair:

```
<name1>=<value1>[&:<name2>=<value2>...&:<nameN>=<valueN>]
```

Configure the following environment variables in the **Cluster Environment Variables** property:

HADOOP_NODE_JDK_HOME

Represents the directory from which you run the cluster services and the JDK version that the cluster nodes use. Required to run the Java transformation in the Hadoop environment and Sqoop mappings on the Blaze engine. Default is /usr/java/default. The JDK version that the Data Integration Service uses must be compatible with the JDK version on the cluster.

Set to <cluster JDK home>/jdk<version>.

For example, HADOOP_NODE_JDK_HOME=<cluster JDK home>/jdk<version>.

Cluster Library Path

Cluster Library Path property is a list of path variables for shared libraries on the cluster. You can add or edit library path variables.

To edit the property in the text box, use the following format with : to separate each path variable:

```
<variable1>[:<variable2>...:<variableN>]
```

Configure the library path variables in the **Cluster Library Path** property.

Common Advanced Properties

Common advanced properties are a list of advanced or custom properties that are unique to the Hadoop environment. The properties are common to the Blaze and Spark engines. Each property contains a name and a value. You can add or edit advanced properties.

To edit the property in the text box, use the following format with `&`: to separate each name-value pair:

```
<name1>=<value1>[&:<name2>=<value2>...&:<nameN>=<valueN>]
```

Configure the following property in the **Advanced Properties** of the common properties section:

infapdo.java.opts

List of Java options to customize the Java run-time environment. The property contains default values.

If mappings in a MapR environment contain a Consolidation transformation or a Match transformation, change the following value:

- `-Xmx512M`. Specifies the maximum size for the Java virtual memory. Default is 512 MB. Increase the value to at least 700 MB.

For example, `infapdo.java.opts=-Xmx700M`

Blaze Engine Advanced Properties

Blaze advanced properties are a list of advanced or custom properties that are unique to the Blaze engine. Each property contains a name and a value. You can add or edit advanced properties.

To edit the property in the text box, use the following format with `&`: to separate each name-value pair:

```
<name1>=<value1>[&:<name2>=<value2>...&:<nameN>=<valueN>]
```

Configure the following properties in the **Advanced Properties** of the Blaze configuration section:

infagrid.cadi.namespace

Namespace for the Data Integration Service to use. Required to set up multiple Blaze instances.

Set to `<unique namespace>`.

For example, `infagrid.cadi.namespace=TestUser1_namespace`

infagrid.blaze.console.jsfport

JSF port for the Blaze engine console. Use a port number that no other cluster processes use. Required to set up multiple Blaze instances.

Set to `<unique JSF port value>`.

For example, `infagrid.blaze.console.jsfport=9090`

infagrid.blaze.console.httpport

HTTP port for the Blaze engine console. Use a port number that no other cluster processes use. Required to set up multiple Blaze instances.

Set to `<unique HTTP port value>`.

For example, `infagrid.blaze.console.httpport=9091`

infagrid.node.local.root.log.dir

Path for the Blaze service logs. Default is `/tmp/infra/logs/blaze`. Required to set up multiple Blaze instances.

Set to `<local Blaze services log directory>`.

For example, `infagrid.node.local.root.log.dir=<directory path>`

infacal.hadoop.logs.directory

Path in HDFS for the persistent Blaze logs. Default is `/var/log/hadoop-yarn/apps/informatica`. Required to set up multiple Blaze instances.

Set to `<persistent log directory path>`.

For example, `infacal.hadoop.logs.directory=<directory path>`

infagrid.node.hadoop.local.root.log.dir

Path in the Hadoop connection for the service log directory.

Set to `<service log directory path>`.

For example, `infagrid.node.local.root.log.dir=$HADOOP_NODE_INFA_HOME/blazeLogs`

Spark Advanced Properties

Spark advanced properties are a list of advanced or custom properties that are unique to the Spark engine. Each property contains a name and a value. You can add or edit advanced properties. Each property contains a name and a value. You can add or edit advanced properties.

Configure the following properties in the **Advanced Properties** of the Spark configuration section:

To edit the property in the text box, use the following format with `&`: to separate each name-value pair:

```
<name1>=<value1>[&:<name2>=<value2>...&:<nameN>=<valueN>]
```

spark.authenticate

Enables authentication for the Spark service on Hadoop. Required for Spark encryption.

Set to `TRUE`.

For example, `spark.authenticate=TRUE`

spark.authenticate.enableSaslEncryption

Enables encrypted communication when SASL authentication is enabled. Required if Spark encryption uses SASL authentication.

Set to `TRUE`.

For example, `spark.authenticate.enableSaslEncryption=TRUE`

spark.executor.cores

Indicates the number of cores that each executor process uses to run tasklets on the Spark engine.

Set to: `spark.executor.cores=1`

spark.executor.instances

Indicates the number of instances that each executor process uses to run tasklets on the Spark engine.

Set to: `spark.executor.instances=1`

spark.executor.memory

Indicates the amount of memory that each executor process uses to run tasklets on the Spark engine.

Set to: `spark.executor.memory=3G`

infaspark.driver.cluster.mode.extraJavaOptions

List of extra Java options for the Spark driver that runs inside the cluster. Required for streaming mappings to read from or write to a Kafka cluster that uses Kerberos authentication.

For example, set to:

```
infaspark.driver.cluster.mode.extraJavaOptions=  
-Djava.security.egd=file:/dev/./urandom  
-XX:MaxMetaspaceSize=256M -Djavax.security.auth.useSubjectCredsOnly=true  
-Djava.security.krb5.conf=<path to keytab file>/krb5.conf  
-Djava.security.auth.login.config=<path to jaas config>/kafka_client_jaas.config
```

To configure the property for a specific user, you can include the following lines of code:

```
infaspark.driver.cluster.mode.extraJavaOptions =  
-Djava.security.egd=file:/dev/./urandom  
-XX:MaxMetaspaceSize=256M -XX:+UseG1GC -XX:MaxGCPauseMillis=500  
-Djava.security.krb5.conf=/etc/krb5.conf
```

infaspark.executor.extraJavaOptions

List of extra Java options for the Spark executor. Required for streaming mappings to read from or write to a Kafka cluster that uses Kerberos authentication.

For example, set to:

```
infaspark.executor.extraJavaOptions=  
-Djava.security.egd=file:/dev/./urandom  
-XX:MaxMetaspaceSize=256M -Djavax.security.auth.useSubjectCredsOnly=true  
-Djava.security.krb5.conf=<path to krb5.conf file>/krb5.conf  
-Djava.security.auth.login.config=<path to JAAS config>/kafka_client_jaas.config
```

To configure the property for a specific user, you can include the following lines of code:

```
infaspark.executor.extraJavaOptions =  
-Djava.security.egd=file:/dev/./urandom  
-XX:MaxMetaspaceSize=256M -XX:+UseG1GC -XX:MaxGCPauseMillis=500  
-Djava.security.krb5.conf=/etc/krb5.conf
```

infaspark.flatfile.writer.nullValue

When the Databricks Spark engine writes to a target, it converts null values to empty strings (" "). For example, 12, AB,"",23p09udj.

The Databricks Spark engine can write the empty strings to string columns, but when it tries to write an empty string to a non-string column, the mapping fails with a type mismatch.

To allow the Databricks Spark engine to convert the empty strings back to null values and write to the target, configure the following advanced property in the Databricks Spark connection:

```
infaspark.flatfile.writer.nullValue=true
```

spark.hadoop.validateOutputSpecs

Validates if the HBase table exists. Required for streaming mappings to write to a HBase target in an Amazon EMR cluster. Set the value to false.

infaspark.json.parser.mode

Specifies the parser how to handle corrupt JSON records. You can set the value to one of the following modes:

- DROPMALFORMED. The parser ignores all corrupted records. Default mode.
- PERMISSIVE. The parser accepts non-standard fields as nulls in corrupted records.
- FAILFAST. The parser generates an exception when it encounters a corrupted record and the Spark application goes down.

infaspark.json.parser.multiLine

Specifies whether the parser can read a multiline record in a JSON file. You can set the value to true or false. Default is false. Applies only to non-native distributions that use Spark version 2.2.x and above.

infaspark.pythontx.exec

Required to run a Python transformation on the Spark engine for Big Data Management. The location of the Python executable binary on the worker nodes in the Hadoop cluster.

For example, set to:

```
infaspark.pythontx.exec=/usr/bin/python3.4
```

If you use the installation of Python on the Data Integration Service machine, set the value to the Python executable binary in the Informatica installation directory on the Data Integration Service machine.

For example, set to:

```
infaspark.pythontx.exec=INFA_HOME/services/shared/spark/python/lib/python3.4
```

infaspark.pythontx.executorEnv.PYTHONHOME

Required to run a Python transformation on the Spark engine for Big Data Management and Big Data Streaming. The location of the Python installation directory on the worker nodes in the Hadoop cluster.

If the Python installation directory on the worker nodes is in a directory such as `usr/lib/python`, set the property to the following value:

```
infaspark.pythontx.executorEnv.PYTHONHOME=usr/lib/python
```

If you use the installation of Python on the Data Integration Service machine, use the location of the Python installation directory on the Data Integration Service machine.

For example, set the property to the following value:

```
infaspark.pythontx.executorEnv.PYTHONHOME=
INFA_HOME/services/shared/spark/python/
```

infaspark.pythontx.executorEnv.LD_PRELOAD

Required to run a Python transformation on the Spark engine for Big Data Streaming. The location of the Python shared library in the Python installation folder on the Data Integration Service machine.

For example, set to:

```
infaspark.pythontx.executorEnv.LD_PRELOAD=
INFA_HOME/services/shared/spark/python/lib/libpython3.6m.so
```

infaspark.pythontx.submit.lib.JEP_HOME

Required to run a Python transformation on the Spark engine for Big Data Streaming. The location of the Jep package in the Python installation folder on the Data Integration Service machine.

For example, set to:

```
infaspark.pythontx.submit.lib.JEP_HOME=
INFA_HOME/services/shared/spark/python/lib/python3.6/site-packages/jep/
```

spark.shuffle.encryption.enabled

Enables encrypted communication when authentication is enabled. Required for Spark encryption.

Set to TRUE.

For example, `spark.shuffle.encryption.enabled=TRUE`

spark.scheduler.maxRegisteredResourcesWaitingTime

The number of milliseconds to wait for resources to register before scheduling a task. Default is 30000. Decrease the value to reduce delays before starting the Spark job execution. Required to improve performance for mappings on the Spark engine.

Set to 15000.

For example, `spark.scheduler.maxRegisteredResourcesWaitingTime=15000`

spark.scheduler.minRegisteredResourcesRatio

The minimum ratio of registered resources to acquire before task scheduling begins. Default is 0.8. Decrease the value to reduce any delay before starting the Spark job execution. Required to improve performance for mappings on the Spark engine.

Set to: 0.5

For example, `spark.scheduler.minRegisteredResourcesRatio=0.5`

APPENDIX B

Multiple Blaze Instances on a Cluster

This appendix includes the following topics:

- [Overview, 127](#)
- [Step 1. Prepare the Hadoop Cluster for the Blaze Engine, 128](#)
- [Step 2. Configure Data Integration Service Properties, 129](#)
- [Step 3. Create a Hadoop Connection, 130](#)
- [Step 4. Configure Additional Hadoop Connection Properties, 131](#)
- [Step 5. Set Mapping Preferences, 132](#)
- [Result, 134](#)

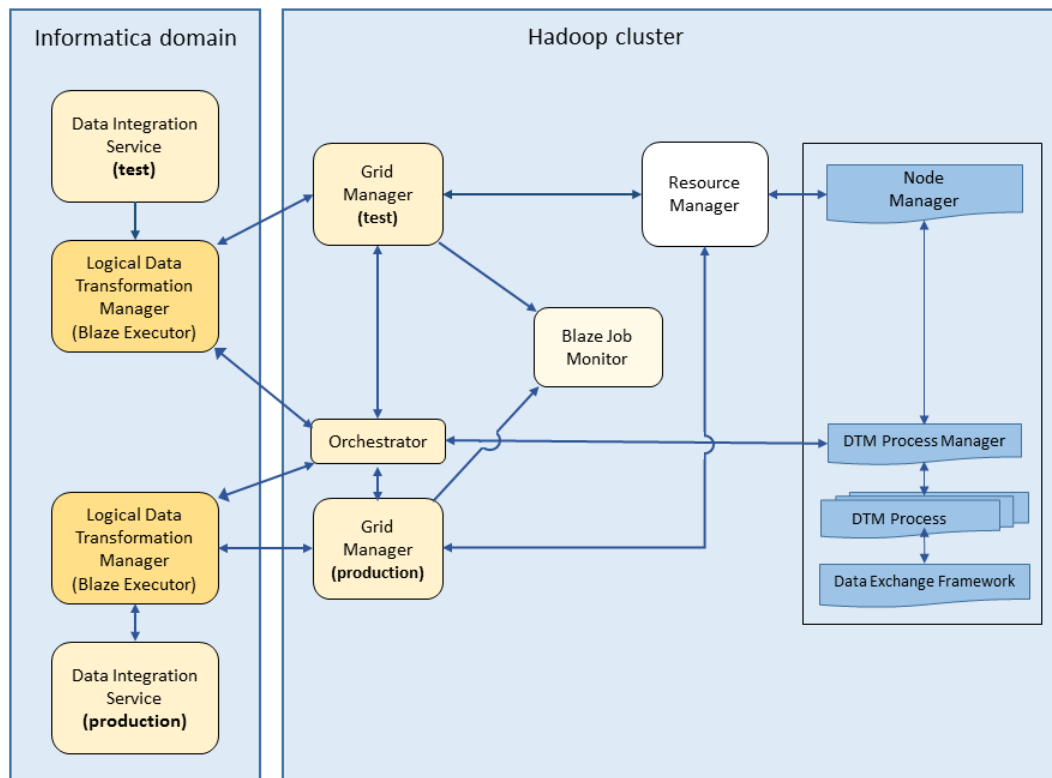
Overview

When you use the Blaze engine to run mappings, Blaze uses a Grid Manager at run time to allot tasks to various nodes in a Hadoop cluster. The Grid Manager aids in resource allocation.

You can use the same Hadoop cluster to stage your test environment and establish a production environment. To control resource use on the cluster, you can establish a separate Blaze instance for testing and another for production.

Each instance requires a separate Grid Manager. You create an additional Grid Manager by performing a series of steps to create separate infrastructure for each Blaze instance, including a unique namespace and a Hadoop connection for each Blaze instance to use. Each Blaze instance also requires a separate Hadoop connection.

The following image shows how a separate Data Integration Service on the domain creates a separate Grid Manager on the cluster:



The image shows how separate Data Integration Services use separate Blaze instances. Each instance uses a separate Grid Manager to communicate with the cluster resource manager to balance resources.

Perform the following steps to set up separate Blaze instances:

- Step 1. Prepare the Hadoop cluster for the Blaze engine.
- Step 2. Configure Data Integration Service properties.
- Step 3. Create a new Hadoop connection.
- Step 4. Configure Additional Hadoop Connection Properties.
- Step 5. Set Mapping Preferences.

Step 1. Prepare the Hadoop Cluster for the Blaze Engine

To run mappings on the Blaze engine, perform the following tasks:

1. Create an account for the Blaze engine user.
2. Create Blaze engine directories and grant permissions.
3. Grant permissions on the Hive source database.

Create a Blaze User Account

On all nodes in the Hadoop cluster, create an operating system user account for the user you want to run the additional Blaze instance. For example, run the following command:

```
useradd testuser1
```

Create Blaze Engine Directories and Grant Permissions

Create the following directories on the Hadoop cluster:

Local services log directory

Configure a local services log directory on all nodes in the cluster and grant permissions to the Blaze user account.

HDFS temporary working directory

Configure a working directory on HDFS for the Blaze engine and grant permissions to the Blaze user account.

Note: This directory is separate from the aggregated persistent log directory.

Verify that a persistent aggregated HDFS log directory exists on the cluster. For example, `/var/log/Hadoop-yarn/apps/Informatica`.

Note: It is not necessary to create a new directory for persistent logs. Both Blaze instances can use the same persistent aggregated HDFS log directory.

Grant Permissions on the Hive Source Database

Grant the Blaze user account CREATE TABLE permission on the Hive source database. The CREATE TABLE permission is required in the following situations:

- The Hive source table uses SQL standard-based authorization.
- A mapping contains a Lookup transformation with an SQL override.

Step 2. Configure Data Integration Service Properties

Configure Data Integration Service properties to enable two Blaze instances on the Hadoop environment.

You can create a Data Integration Service, or configure one that has not run mappings using the Blaze engine.

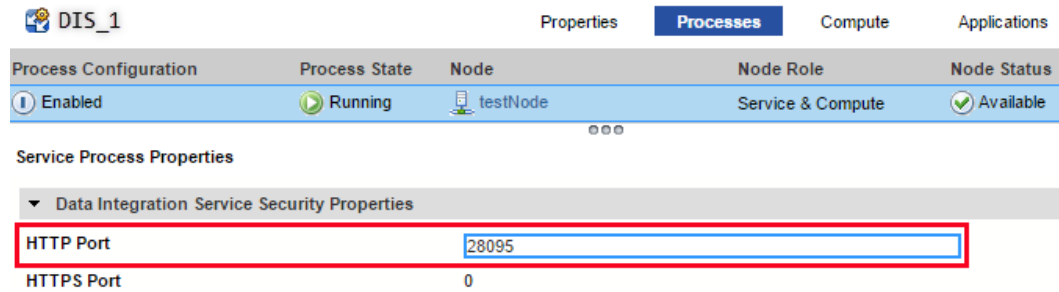
Configure Data Integration Service properties in the Administrator tool.

Configure Data Integration Service Process Properties

Configure the HTTP port property on the **Processes** tab:

Property	Description
HTTP Port	The port that the Data Integration Service uses to communicate with the cluster over HTTP. Configure the port with a number that no other process uses.

The following image shows the HTTP Port property:



Step 3. Create a Hadoop Connection

Create a Hadoop connection for the Blaze instance to use.

1. In the **Connections** tab of the Administrator tool, right-click the domain node and choose **New > Connection**. Choose the connection type **Cluster > Hadoop**.
The New Connections wizard opens.
2. In Step 1 of the New Connection wizard, select the cluster configuration for the Hadoop connection to use.
3. Configure the Impersonation User Name property with the same impersonation user that you configured in Step 1, "Create a Blaze User Account."
4. Step 2 of the New Connection wizard requires no mandatory input. Click **Next**.
When you click through this screen without configuring properties, the values are populated from the cluster configuration.
5. In Step 3 of the New Connection wizard, configure the Blaze Staging Directory property with the path that you configured on the cluster in "Create Blaze Engine Directories and Grant Permissions."
6. Configure the Blaze User Name property with the same user name that you used to configure the Impersonation User in Step 1 of this topic.
7. Configure the Minimum Port and Maximum Port properties with a port range for the connection.
You can supply a range within the port range of an existing Grid Manager, as long as ports are available when the mapping runs. The default range is 300 ports.
8. Click **Finish** to create the new connection.

Step 4. Configure Additional Hadoop Connection Properties

Configure a Hadoop connection for each Blaze instance.

Note: Prior to version 10.2.1, the properties on this page were configured in the `hadoopEnv.properties` file. If you want to use the property values from that file, back up the file before you upgrade. Then configure the values in the environment variables property in the Hadoop connection.

Optionally Create a New Namespace

When the machine where the Data Integration Service runs contains two domains running on the same version of Informatica, you configure a new Blaze instance on the domain where you want to run the new Blaze instance.

Configure the following property in the Hadoop connection Blaze Advanced Properties:

infagrid.cadi.namespace

Namespace for the Data Integration Service to use.

Configure the property as follows:

```
infagrid.cadi.namespace=<unique value>
```

For example,

```
infagrid.cadi.namespace=TestUser1_namespace
```

Configure the Blaze Job Monitor Address

Configure the following property in the Hadoop connection Blaze Advanced Properties:

Blaze Job Monitor Address

The host name and port number for the Blaze Job Monitor.

Use the following format:

```
<hostname>:<port>
```

Where

- `<hostname>` is the host name or IP address of the Blaze Job Monitor server.
- `<port>` is the port on which the Blaze Job Monitor listens for remote procedure calls (RPC).

For example, enter: `myhostname:9080`

Configure Ports

Configure the following property in the Hadoop connection Blaze Advanced Properties. Use port numbers that no other cluster processes use.

infagrid.blaze.console.jsfport

JSF port for the Blaze engine console.

Configure the property as follows:

```
infagrid.blaze.console.jsfport=<unique value>
```

For example,

```
infagrid.blaze.console.jsfport=9090
```

Configure Directory Paths

To configure the following properties in Blaze Advanced Properties, click the **Edit** icon in the Blaze Configuration section of Hadoop connection properties, and then click the **Edit** icon adjacent to the Advanced properties pane.

The following image shows the Edit icon for Advanced properties:

infagrid.node.local.root.log.dir

Path for the Blaze service logs.

Note: This is the path that you configured in Step 1 as the local services log directory.

Configure the property as follows:

```
infagrid.node.local.root.log.dir=<directory path>
```

Default:

```
infagrid.node.local.root.log.dir=/tmp/infra/logs/blaze
```

infacal.hadoop.logs.directory

Path in HDFS for the persistent Blaze logs.

Note: This is the path that you configured in Step 1 as the persistent log directory.

Configure the property as follows:

```
infacal.hadoop.logs.directory=<directory path>
```

Default:

```
infacal.hadoop.logs.directory=infacal.hadoop.logs.directory=/var/log/hadoop-yarn/apps/informatica
```

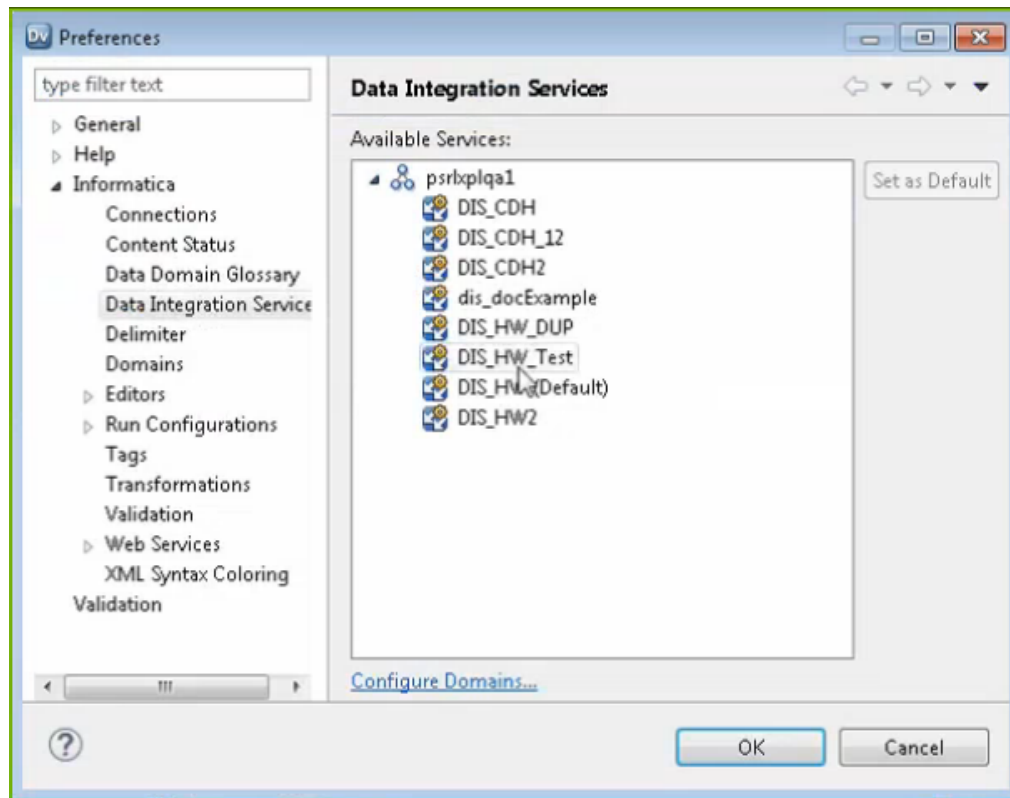
Step 5. Set Mapping Preferences

Before you run the mapping in the Developer tool, configure the mapping to use the Data Integration Service and Hadoop connection you want to use to run the mapping.

1. In the Developer tool, select **Mapping > Preferences**.

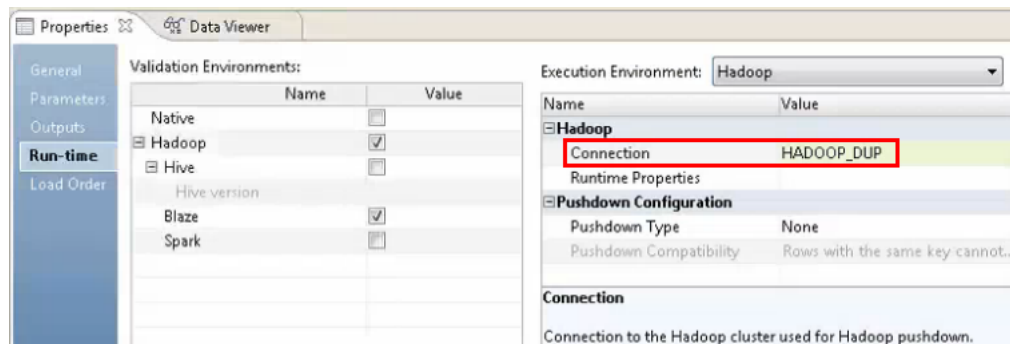
- Expand the **Informatica** node, and then select **Data Integration Service**.

The following image shows the list of available services in the **Preferences** window:



- Select the Data Integration Service that you want to use, and then click **OK**.
- In the **Properties** tab of the mapping, select the **Run-time** sub-tab.
- In the Execution Environment area, set the Connection property to the Hadoop connection that you created.

The following image shows the Connection property in the **Properties** tab:



Result

The Data Integration Service creates a Grid Manager on the cluster the first time that it runs a mapping using the Blaze engine.

After you run the mapping, you can verify that the mapping used the Data Integration Service and new Grid Manager that you intended to use to run the mapping. Verify the resources that the mapping used by examining the Running Applications list in the Hadoop Resource Manager web interface. Look for applications that correspond to the namespace that you configured for the Blaze instance.

The following image shows applications with a name that includes the namespace, "testuser1_namespace," that you configured for the Grid Manager:

The screenshot shows the Hadoop Resource Manager web interface. On the left is a navigation menu with options like 'Cluster', 'About Nodes', 'Applications', and 'Scheduler'. The main area is titled 'RUNNING Applications' and displays 'Cluster Metrics' and 'User metrics for testUser1'. Below these are two tables showing application details.

Cluster Metrics										
Apps Submitted	Apps Pending	Apps Running	Apps Completed	Containers Running	Memory Used	Memory Total	Memory Reserved	VCores Used	VCores Total	
63	0	2	61	2	3 GB	16 GB	0 B	2	12	

User metrics for testUser1									
Apps Submitted	Apps Pending	Apps Running	Apps Completed	Containers Running	Containers Pending	Containers Reserved	Memor Used		
0	0	0	0	0	0	0	0 B		

ID	User	Name	Application Type	Queue	StartTime	FinishTime	State
application_1481793531093_0063	devbld	/testUser1_namespace /Blaze Grid Manager	Informatica Blaze	root.users.devbld	Wed Jan 25 11:10:27 -0800 2017	N/A	RUNNING
application_1481793531093_0062	devbld	/testUser1_namespace /Blaze Job Monitor	Informatica Blaze	root.users.devbld	Wed Jan 25 11:09:28 -0800 2017	N/A	RUNNING

After completion of the mapping run, the Grid Manager persists. The mapping uses the same Grid Manager whenever it runs with the unique combination of Data Integration Service and connection.

To use multiple connections that use the same Grid Manager, use an identical namespace in each connection to refer to the Blaze instance. Verify that each connection also uses identical values for the Blaze user name and queue name. If you use different values for the Blaze user name and queue name to connect to the same Blaze instance, the mapping fails.

APPENDIX C

Monitoring REST API

This appendix includes the following topics:

- [Monitoring REST API Overview, 135](#)
- [Monitoring Metadata Document, 136](#)
- [MappingStats, 138](#)
- [MappingAdvancedStats, 140](#)
- [MappingExecutionSteps, 145](#)
- [MappingExecutionPlans, 149](#)

Monitoring REST API Overview

You can monitor statistics retrieved through the REST API.

With the REST API, you can get the mapping execution statistics for Big Data Management mappings that run on the Data Integration Service or the Blaze or Spark engines. You can get mapping statistics, mapping advanced statistics, mapping execution steps, and mapping execution plans.

Note: Before you use the REST APIs, identify the client or the programming language that you want to use to call the REST APIs. Ensure that you start the REST Operations Hub Service to use the REST APIs. Enable the monitoring Model Repository Service to get data from the REST APIs.

You need to enter the mapping jobID as input to the REST API to get the mapping execution statistics. Monitoring REST API requires user name, password encrypted with the `pmpasswd` command line program, and security domain to connect to the Informatica domain to get the monitoring execution statistics. You can pass the Informatica domain user name, password, and security domain as headers during the REST API.

You can use the following REST APIs to get the mapping execution statistics:

- MappingStats
- MappingAdvancedStats
- MappingExecutionSteps
- MappingExecutionPlans

Monitoring Metadata Document

Use the metadata document to get a list of monitoring REST APIs. With the monitoring REST APIs, you can get mapping execution statistics and the input and output parameters for each REST API.

You can use the following URL to view the metadata document:

```
<RESTOperationsHubService_Host>:<RESTOperationsHubService_Port>/RestOperationsHub/  
services/v1/MappingService/$metadata
```

Sample Metadata Document

The following sample metadata document displays the mapping execution information for each REST API:

```
<?xml version='1.0' encoding='UTF-8'?>  
<edmx:Edmx Version="4.0" xmlns:edmx="http://docs.oasis-open.org/odata/ns/edmx">  
  <edmx:DataServices>  
    <Schema xmlns="http://docs.oasis-open.org/odata/ns/edm" Namespace="OData.Domain"  
      Alias="MappingService">  
      <EntityType Name="MappingAdvancedStat">  
        <Key>  
          <PropertyRef Name="jobId"/>  
        </Key>  
        <Property Name="jobId" Type="Edm.String"/>  
        <Property Name="mappingStat" Type="MappingService.MSummaryStatDetails"/>  
        <Property Name="detailedStats" Type="MappingService.MDetailedStats"/>  
        <Property Name="summaryStats" Type="MappingService.MSrcTgtStats"/>  
      </EntityType>  
      <EntityType Name="MappingExecutionPlan">  
        <Key>  
          <PropertyRef Name="jobId"/>  
        </Key>  
        <Property Name="jobId" Type="Edm.String"/>  
        <Property Name="scriptList"  
Type="Collection(MappingService.ScriptEntryStat)"/>  
      </EntityType>  
      <EntityType Name="MappingStat">  
        <Key>  
          <PropertyRef Name="jobId"/>  
        </Key>  
        <Property Name="jobId" Type="Edm.String"/>  
        <Property Name="mappingDetails"  
Type="MappingService.MSummaryStatDetails"/>  
      </EntityType>  
      <EntityType Name="MappingExecutionStep">  
        <Key>  
          <PropertyRef Name="jobId"/>  
          <PropertyRef Name="fetchScripts"/>  
        </Key>  
        <Property Name="jobId" Type="Edm.String"/>  
        <Property Name="fetchScripts" Type="Edm.Boolean"/>  
        <Property Name="executionStepStats"  
Type="Collection(MappingService.MExecutionStepStat)"/>  
      </EntityType>  
      <ComplexType Name="MHiveScriptStat"  
BaseType="MappingService.MExecutionStepStat">  
        <Property Name="hiveHistoryFile" Type="Edm.String"/>  
        <Property Name="hiveSessionFile" Type="Edm.String"/>  
      </ComplexType>  
      <ComplexType Name="MGroupStat">  
        <Property Name="processedBytes" Type="Edm.Int64"/>  
        <Property Name="processedRows" Type="Edm.Int64"/>  
        <Property Name="bytesThroughput" Type="Edm.Int64"/>  
        <Property Name="rowsThroughput" Type="Edm.Int64"/>  
        <Property Name="errorRows" Type="Edm.Int64"/>  
        <Property Name="errorBytes" Type="Edm.Int64"/>  
        <Property Name="groupName" Type="Edm.String"/>  
      </ComplexType>  
    </Schema>  
  </edmx:DataServices>  
</edmx:Edmx>
```



```

        <Property Name="firstRowTime" Type="Edm.Int64"/>
    </ComplexType>
    <ComplexType Name="MDetailedStats">
        <Property Name="memoryData" Type="Collection(Edm.Int64)"/>
        <Property Name="createdTime" Type="Collection(Edm.Int64)"/>
        <Property Name="cpuData" Type="Collection(Edm.Int64)"/>
        <Property Name="sourceTargetStats"
Type="Collection(MappingService.TxInstanceStat)"/>
        <Property Name="lastPurgeTime" Type="Edm.Int64"/>
    </ComplexType>
    <ComplexType Name="MNameValuePair">
        <Property Name="name" Type="Edm.String"/>
        <Property Name="value" Type="Edm.String"/>
    </ComplexType>
    <ComplexType Name="MExecutionStepStat">
        <Property Name="endTime" Type="Edm.Int64"/>
        <Property Name="startTime" Type="Edm.Int64"/>
        <Property Name="errorCode" Type="Edm.String"/>
        <Property Name="errorMessage" Type="Edm.String"/>
        <Property Name="name" Type="Edm.String"/>
        <Property Name="state" Type="Edm.String"/>
        <Property Name="subExecutionStepStat"
Type="Collection(MappingService.MExecutionStepStat)"/>
    </ComplexType>
    <ComplexType Name="MLocalMappingStepStat"
BaseType="MappingService.MExecutionStepStat"/>
    <ComplexType Name="MGridTaskStepStat"
BaseType="MappingService.MExecutionStepStat">
        <Property Name="percentageCompletion" Type="Edm.Double"/>
        <Property Name="webPageURL" Type="Edm.String"/>
        <Property Name="incomingDependencies" Type="Edm.String"/>
        <Property Name="outgoingDependencies" Type="Edm.String"/>
    </ComplexType>
    <ComplexType Name="TxInstanceStat">
        <Property Name="instanceName" Type="Edm.String"/>
        <Property Name="isSource" Type="Edm.Boolean"/>
        <Property Name="bytes" Type="Collection(Edm.Int64)"/>
        <Property Name="rows" Type="Collection(Edm.Int64)"/>
    </ComplexType>
    <ComplexType Name="ScriptEntryStat">
        <Property Name="name" Type="Edm.String"/>
        <Property Name="content" Type="Edm.String"/>
        <Property Name="depends" Type="Edm.String"/>
    </ComplexType>
    <ComplexType Name="MCommandTaskStat"
BaseType="MappingService.MExecutionStepStat"/>
    <ComplexType Name="MSrcTgtStats">
        <Property Name="processStatSummary"
Type="MappingService.MProcessStatSummary"/>
        <Property Name="sourceTxStats"
Type="Collection(MappingService.MSourceTxStat)"/>
        <Property Name="targetTxStats"
Type="Collection(MappingService.MTargetTxStat)"/>
    </ComplexType>
    <ComplexType Name="MTargetTxStat" BaseType="MappingService.AbstractTxStat"/>
    <ComplexType Name="MSparkApplicationStepStat"
BaseType="MappingService.MExecutionStepStat">
        <Property Name="jobTrackerUrl" Type="Edm.String"/>
        <Property Name="query" Type="Edm.String"/>
    </ComplexType>
    <ComplexType Name="MHadoopTezJobStat"
BaseType="MappingService.MExecutionStepStat">
        <Property Name="dagProgress" Type="Edm.Double"/>
        <Property Name="hadoopJobPageUrl" Type="Edm.String"/>
        <Property Name="hadoopOfficialTezJobId" Type="Edm.String"/>
    </ComplexType>
    <ComplexType Name="MSegmentStepStat"
BaseType="MappingService.MExecutionStepStat">
        <Property Name="numOfFailedTasklets" Type="Edm.Int64"/>
        <Property Name="numOfSucceededTasklets" Type="Edm.Int64"/>
        <Property Name="numOfTasklets" Type="Edm.Int64"/>

```

```

        <Property Name="percentageCompletion" Type="Edm.Double"/>
        <Property Name="webPageURL" Type="Edm.String"/>
    </ComplexType>
    <ComplexType Name="MHiveQueryStat"
BaseType="MappingService.MExecutionStepStat">
        <Property Name="hiveOfficalHiveQueryId" Type="Edm.String"/>
        <Property Name="query" Type="Edm.String"/>
    </ComplexType>
    <ComplexType Name="MProcessStatSummary">
        <Property Name="avgCpuUsage" Type="Edm.Double"/>
        <Property Name="avgMemUsage" Type="Edm.Int64"/>
    </ComplexType>
    <ComplexType Name="AbstractTxStat">
        <Property Name="instanceName" Type="Edm.String"/>
        <Property Name="groupStats"
Type="Collection (MappingService.MGroupStat)"/>
    </ComplexType>
    <ComplexType Name="MSummaryStatDetails">
        <Property Name="status" Type="Edm.String"/>
        <Property Name="mappingName" Type="Edm.String"/>
        <Property Name="applicationName" Type="Edm.String"/>
        <Property Name="serviceName" Type="Edm.String"/>
        <Property Name="logFileName" Type="Edm.String"/>
        <Property Name="startTime" Type="Edm.Int64"/>
        <Property Name="endTime" Type="Edm.Int64"/>
        <Property Name="executorType" Type="Edm.String"/>
        <Property Name="executingNode" Type="Edm.String"/>
        <Property Name="userName" Type="Edm.String"/>
        <Property Name="securityDomain" Type="Edm.String"/>
    </ComplexType>
    <ComplexType Name="MUserDefinedTaskStat"
BaseType="MappingService.MExecutionStepStat">
        <Property Name="statsProperties"
Type="Collection (MappingService.MNameValuePair)"/>
    </ComplexType>
    <ComplexType Name="MSourceTxStat" BaseType="MappingService.AbstractTxStat"/>
    <ComplexType Name="MHadoopMapReduceJobStat"
BaseType="MappingService.MExecutionStepStat">
        <Property Name="hadoopJobPageUrl" Type="Edm.String"/>
        <Property Name="hadoopOfficialMapReduceJobId" Type="Edm.String"/>
        <Property Name="mapperProgress" Type="Edm.Double"/>
        <Property Name="reducerProgress" Type="Edm.Double"/>
    </ComplexType>
    <EntityContainer Name="MSCONTAINER">
        <EntitySet Name="MappingAdvancedStats"
EntityType="MappingService.MappingAdvancedStat"/>
        <EntitySet Name="MappingExecutionPlans"
EntityType="MappingService.MappingExecutionPlan"/>
        <EntitySet Name="MappingStats" EntityType="MappingService.MappingStat"/>
        <EntitySet Name="MappingExecutionSteps"
EntityType="MappingService.MappingExecutionStep"/>
    </EntityContainer>
</Schema>
</edmx:DataServices>
</edmx:Edmx>

```

MappingStats

With the MappingStats REST API, you can view basic mapping statistics, such as the name and time taken for the job for mappings that run in the native or Hadoop environment.

GET Request

To request information from the MappingStats, use the following URL:

<RESTOperationsHubService_Host>:<RESTOperationsHubService_Port>/RestOperationsHub/services/v1/MappingService/MappingStats('jobId')

The following table describes the attributes in the MappingStats Get URL:

Field	Type	Description
userName	String	Required. User name to connect to the domain. You can pass the input value as a header.
encryptedpassword	String	Required. Password for the user. Encrypt the password with the pmpasswd command line program. You can pass the input value as a header.
securityDomain	String	Optional. The security domain to which the domain user belongs. You can pass the input value as a header. Default is Native.
jobId	String	Required. The argument of the entity that contains the ID for the mappings. You can pass the property as part of the URL.

Get Response

Return information for the MappingStats for the specified Job ID.

The following table describes the MappingStats attributes present in the body of the response on the native or Hadoop environment:

Field	Type	Description
jobId	String	The argument of the entity that contains the ID for the mapping.
mappingDetails	n/a	Container for mapping run information, such as mapping name, start time, and end time.
status	String	State of the job run.
mappingName	String	Name of the mapping.
applicationName	String	Name of the application.
serviceName	String	Name of the service.
logFileName	String	Location of the log file.
startTime	Integer	Start time of the job.
endTime	Integer	End time of the job.
executorType	String	Type of the run-time environment where you run the mapping.
executingNode	String	The Data Integration Service node where the job ran.

Field	Type	Description
userName	String	User who started the job.
securityDomain	String	The security domain to which the domain user belongs.

Sample Retrieve Mapping Statistics

The sample use case is to use the script to retrieve the details of the mapping statistics on the Spark environment.

You can use the REST API to retrieve information about the mapping statistics with the following request URL for a mapping with Job ID as `_TNo09ELEeimY76kFyfuw`:

```
<RESTOperationsHubService_Host>:<RESTOperationsHubService_Port>/RestOperationsHub/
services/v1/MappingService/MappingStats('_TNo09ELEeimY76kFyfuw')
```

Mapping Statistics Output

```
{
  "@odata.context": "$metadata#MappingStats/$entity",
  "jobId": "_TNo09ELEeimY76kFyfuw",
  "mappingDetails": {
    "status": "COMPLETED",
    "mappingName": "HDFSTgtAppend_MultiPartition_SparkMode",
    "applicationName": "HDFSTargetAppend",
    "serviceName": "DIS_HDP_2.6",
    "logFileName": "/data/Informatica/10.2.2_252/logs/node252/services/
DataIntegrationService/disLogs/ms/
DEPLOYED_MAPPING_HDFSTargetAppend_HDFSTgtAppend_MultiPartition_SparkMode-
_TNo09ELEeimY76kFyfuw_20181016_115325_006.log",
    "startTime": 1539671005830,
    "endTime": 1539671244752,
    "executorType": "SPARK",
    "executingNode": "node252",
    "userName": "Administrator",
    "securityDomain": "Native"
  }
}
```

MappingAdvancedStats

With the MappingAdvancedStats REST API, you can view mapping summary and detailed statistics, such as the job throughput for mappings that run in the native or Hadoop environment.

GET Request

To request information from the MappingAdvancedStats, use the following URL:

```
<RESTOperationsHubService_Host>:<RESTOperationsHubService_Port>/RestOperationsHub/
services/v1/MappingService/MappingAdvancedStats('jobId')
```

The following table describes the attributes in the MappingAdvancedStats Get URL:

Field	Type	Description
userName	String	Required. User name to connect to the domain. You can pass the input value as a header.
encryptedpassword	String	Required. Password for the user. Encrypt the password with the pmpasswd command line program. You can pass the input value as a header.
securityDomain	String	Optional. The security domain to which the domain user belongs. You can pass the input value as a header. Default is Native.
jobId	String	Required. The argument of the entity that contains the ID for the mappings. You can pass the property as part of the URL.

Get Response

Return information for the MappingAdvancedStats for the specified Job ID.

The following table describes the MappingAdvancedStats attributes present in the body of the response for native or Hadoop environment:

Field	Type	Description
jobId	String	The argument of the entity that contains the ID for the mapping.
mappingStat	n/a	Container for mapping run information, such as mapping name, start time, and end time.
status	String	State of the job run.
mappingName	String	Name of the mapping.
applicationName	String	Name of the application.
serviceName	String	Name of the service.
logFileName	String	Location of the log file.
startTime	Integer	Start time of the job.
endTime	Integer	End time of the job.
executorType	String	Type of the run-time environment where you run the mapping.
executingNode	String	The Data Integration Service node where the job ran.
userName	String	User who started the job.
securityDomain	String	The security domain to which the domain user belongs.
detailedStats	n/a	Container for detailed mapping statistics.

Field	Type	Description
memoryData	String	Metrics for memory.
createdTime	Integer	Time taken to created the job.
cpuData	Integer	Data for CPU.
sourceTargetStats	n/a	Container for source and target statistics.
instanceName	String	Object name for source or target.
isSource	Integer	For the source, specify isSource as true. For the target, specify isSource as false.
bytes	Integer	Average number of bytes read per second for source or target.
rows	Integer	Number of rows read for source or target.
lastPurgeTime	Integer	Last time of the purge task.
summaryStats	n/a	Container for the summary statistics.
processStatSummary	n/a	Container for the CPU and memory usage.
avgCpuUsage	Integer	Average CPU usage.
avgMemUsage	Integer	Average memory usage.
sourceTxStats	n/a	Container for the source statistics.
groupStats	n/a	Container for processed and throughput statistics.
ProcessedBytes	Integer	Bytes processed for source or target.
ProcessedRows	Integer	Rows processed for source or target.
bytesThroughput	Integer	Throughput in bytes for source or target.
rowsThroughput	Integer	Throughput of rows for source or target.
errorRows	Integer	Error in rows for source or target.
errorBytes	Integer	Error in bytes for source or target.
groupName	String	Name of the group for source or target.
firstRowTime	Integer	Time taken for the first row for source or target.
targetTxStats	n/a	Container for target statistics.

Sample Retrieve Advanced Mapping Statistics

The sample use case is to use the script to retrieve the details of the advanced mapping statistics on the Spark environment.

You can use the REST API to retrieve information about the advanced mapping statistics with the following request URL for a mapping with Job ID as `_TNo09ELEeimY76kFyfuw`:

```
<RESTOperationsHubService_Host>:<RESTOperationsHubService_Port>/RestOperationsHub/
services/v1/MappingService/MappingAdvancedStats('_TNo09ELEeimY76kFyfuw')
```

Advanced Mapping Statistics Output

```
{
  "@odata.context": "$metadata#MappingAdvancedStats/$entity",
  "jobId": "_TNo09ELEeimY76kFyfuw",
  "mappingStat": {
    "status": "COMPLETED",
    "mappingName": "HDFSTgtAppend_MultiPartition_SparkMode",
    "applicationName": "HDFSTargetAppend",
    "serviceName": "DIS_HDP_2.6",
    "logFileName": "/data/Informatica/10.2.2_252/logs/node252/services/
DataIntegrationService/disLogs/ms/
DEPLOYED_MAPPING_HDFSTargetAppend_HDFSTgtAppend_MultiPartition_SparkMode-
TNo09ELEeimY76kFyfuw_20181016_115325_006.log",
    "startTime": 1539671005830,
    "endTime": 1539671244752,
    "executorType": "SPARK",
    "executingNode": "node252",
    "userName": "Administrator",
    "securityDomain": "Native"
  },
  "detailedStats": {
    "memoryData": [],
    "createdTime": [
      1539671058384
    ],
    "cpuData": [],
    "sourceTargetStats": [
      {
        "instanceName": "Read_students_5",
        "isSource": true,
        "bytes": [
          -1
        ],
        "rows": [
          10
        ]
      },
      {
        "instanceName": "Read_students_HDFS_src",
        "isSource": true,
        "bytes": [
          -1
        ],
        "rows": [
          10
        ]
      },
      {
        "instanceName": "Read_student",
        "isSource": true,
        "bytes": [
          -1
        ],
        "rows": [
          10
        ]
      },
      {
        "instanceName": "Write_HDFSAppendTarget",
        "isSource": false,
        "bytes": [
          -1
        ],
        "rows": [
          28
        ]
      }
    ]
  }
}
```

```

    ]
  },
  "lastPurgeTime": 0
},
"summaryStats": {
  "processStatSummary": {
    "avgCpuUsage": 0,
    "avgMemUsage": 0
  },
  "sourceTxStats": [
    {
      "instanceName": "Read_students_5",
      "groupStats": [
        {
          "processedBytes": -1,
          "processedRows": 10,
          "bytesThroughput": -1,
          "rowsThroughput": 10,
          "errorRows": -1,
          "errorBytes": -1,
          "groupName": "Read_students_5",
          "firstRowTime": 0
        }
      ]
    },
    {
      "instanceName": "Read_student",
      "groupStats": [
        {
          "processedBytes": -1,
          "processedRows": 10,
          "bytesThroughput": -1,
          "rowsThroughput": 10,
          "errorRows": -1,
          "errorBytes": -1,
          "groupName": "Read_student",
          "firstRowTime": 0
        }
      ]
    },
    {
      "instanceName": "Read_students_HDFS_src",
      "groupStats": [
        {
          "processedBytes": -1,
          "processedRows": 10,
          "bytesThroughput": -1,
          "rowsThroughput": 5,
          "errorRows": -1,
          "errorBytes": -1,
          "groupName": "Read_students_HDFS_src",
          "firstRowTime": 0
        }
      ]
    }
  ],
  "targetTxStats": [
    {
      "instanceName": "Write_HDFSAppendTarget",
      "groupStats": [
        {
          "processedBytes": -1,
          "processedRows": 28,
          "bytesThroughput": -1,
          "rowsThroughput": 14,
          "errorRows": -1,
          "errorBytes": -1,
          "groupName": "Write_HDFSAppendTarget",
          "firstRowTime": 0
        }
      ]
    }
  ]
}

```



```

    }
  }
}

```

MappingExecutionSteps

With the MappingExecutionStats REST API, you can view execution steps for Hadoop jobs. You can also view the scripts for Hadoop jobs with the fetchScripts option.

GET Request

To request information from the MappingExecutionSteps with the associated scripts within the Hadoop scripts, use the following URL:

```

<RESTOperationsHubService_Host>:<RESTOperationsHubService_Port>/RestOperationsHub/
services/v1/MappingService/MappingExecutionSteps(jobId='id',fetchScripts=true)

```

To request information from the MappingExecutionSteps without the associated scripts within the Hadoop scripts, use the following URL:

```

<RESTOperationsHubService_Host>:<RESTOperationsHubService_Port>/RestOperationsHub/
services/v1/MappingService/MappingExecutionSteps(jobId='id',fetchScripts=false)

```

The following table describes the required attributes in the MappingExecutionSteps Get URL:

Field	Type	Description
userName	String	Required. User name to connect to the domain. You can pass the input value as a header.
encryptedpassword	String	Required. Password for the user. Encrypt the password with the pmpasswd command line program. You can pass the input value as a header.
securityDomain	String	Optional. The security domain to which the domain user belongs. You can pass the input value as a header.
jobId	String	Required. The argument of the entity that contains the ID for the mappings. You can pass the input value as a query.
fetchScripts	String	Required. Gets all associated scripts within the Hadoop scripts. You can enter true to get the associated run tasks with the scripts. To get the associated run tasks without the scripts, enter false.

Get Response

Return information for the MappingExecutionSteps for the specified Job ID.

The following table describes the MappingExecutionSteps attributes present in the body of the response for Blaze environment:

Field	Description
Session Task	Response can include three types of session tasks that are pre session task, main session task, and post session task. The main session task can contain Submappings.
Instances	Transformation or object name.
endTime	End time of the job.
startTime	Start time of the job.
errorCode	Error code.
name	Name of the submapping task based on the run-time environment.
state	State of the mapping run.
subExecutionStepStat	Statistics for the sub execution step.
numOfFailedTasklets	Number of failed tasklets.
numOfSucceededTasklets	Number of succeeded tasklets.
numOfTasklets	Number of tasklets.
percentageCompletion	Percentage completion of tasklets.
webPageURL	Blaze monitoring tasklet URL.
incomingDependencies	Incoming dependencies contain pre session tasks.
outgoingDependencies	Outgoing dependencies contain post session tasks.

The following table describes the MappingExecutionSteps attributes present in the body of the response for Spark environment:

Field	Type	Description
jobId	String	The argument of the entity that contains the ID for the mapping.
fetchScripts	String	Gets all associated scripts within the Spark scripts. You can enter true to get the associated run tasks with the scripts. To get the associated run tasks without the scripts, enter false.
executionStepStats	n/a	Container for the statistics related to mapping execution steps.
endTime	Integer	End time of the job.
startTime	Integer	Start time of the job.
errorCode	Integer	Error code.

Field	Type	Description
errorMessage	Integer	Error message.
name	String	Name of the main mapping task based on the run-time environment.
state	String	State of the job run.
subExecutionStepStats	Integer	Statistics for the subexecution step statistics.
jobTrackerURL	String	The URL to the job tracker.
query	String	The main Spark query.

Sample Retrieve Mapping Execution Steps

The sample use case is to use the script to retrieve the details of the mapping execution steps for the Spark environment.

You can use the REST API to retrieve information about the mapping execution steps with the following request URL for a mapping with Job ID as `_TNo09ELEeiimY76kFyfuw`:

```
<RESTOperationsHubService_Host>:<RESTOperationsHubService_Port>/RestOperationsHub/
services/v1/MappingService/
MappingExecutionSteps(jobId='_TNo09ELEeiimY76kFyfuw',fetchScripts=true)
```

Mapping Execution Steps Output

```
{
  "@odata.context": "$metadata#MappingExecutionSteps/$entity",
  "jobId": "_TNo09ELEeiimY76kFyfuw",
  "fetchScripts": true,
  "executionStepStats": [
    {
      "@odata.type": "#OData.Domain.MSparkApplicationStepStat",
      "endTime": 1539671244359,
      "startTime": 1539671066452,
      "errorCode": "",
      "errorMessage": "",
      "name": "InfaSpark0",
      "state": "COMPLETED",
      "subExecutionStepStat": [],
      "jobTrackerUrl":
        "https://ivlhd621.informatica.com:8090/proxy/application_1539263790092_1418",
      "query": "package com.informatica.exec\n\nimport
com.informatica.bootstrap.functions._\nimport com.informatica.bootstrap.InfaParams._
\nimport com.informatica.bootstrap.InfaStreaming.writeToKafka\nimport
com.informatica.products.infatransform.spark.boot._\nimport com.informatica.bootstrap._
\nimport com.informatica.hive._\nimport com.informatica.bootstrap.{JsonProtocol => JP}
\nimport org.apache.spark._\nimport org.apache.spark.rdd._\nimport
org.apache.spark.storage.StorageLevel._\nimport org.apache.spark.sql._\nimport
org.apache.spark.sql.types._\nimport org.apache.spark.sql.functions._\nimport
org.apache.spark.sql.functions.{ broadcast => infabroadcast }\nimport
org.apache.spark.sql.infa.expressions._\nimport java.io._\nimport java.sql.Timestamp
\nimport scala.reflect.ClassTag\nimport org.apache.spark.sql.catalyst.expressions.Caster
\nimport org.apache.spark.sql.catalyst.expressions.JavaCaster\nimport
com.informatica.bootstrap.JavaTx._\nimport org.apache.spark.Accumulator\nimport
org.apache.spark.util.LongAccumulator\nimport org.apache.spark.scheduler.SparkListener
\nimport org.apache.spark.SparkEnv\nimport org.apache.spark.sql.Row\n\nobject InfaSpark0
{\n  def main(s:Array[String]) {\n    val sc = SparkContextLoader.getSparkContext\n
val sqlContext = SparkContextLoader.getSQLContext\n    val ls = new
LiveStream(sc.getConf)\n    ls.relay(JP.sparkConfToJson(sc.getConf)) \n
ls.relay(JP.hadoopConfToJson(sc.hadoopConfiguration)) \n    val lis = new
InfaListener(ls,\"TAG\")\n    sc.addSparkListener(lis)
```

```

\nsqlContext.sparkSession.experimental.extraPreprocessing = new InfaTaggingRules().rules
\n val accs = List()\n ls.relay(JP.sparkAppDetailsToJson(sc.getConf, accs)) \n
lis.accumulators = accs\n import sqlContext.implicit._\n import
org.apache.spark.sql.functions.{stddev_samp, var_samp}\n val icast = caster("\MM/DD/
YYYY HH24:MI:SS")\n val acast = adapterCaster()\n val jcast = JavaCaster()\n\n
try {\n Tuple2(sqlContext.sql(InfaParams.resolve(InfaParams[String]("\InfaSpark0",
\n0\n))), sqlContext.sql(InfaParams.resolve(InfaParams[String]("\InfaSpark0", \n1\n))));
\n Tuple2(sqlContext.sql(InfaParams.resolve(InfaParams[String]("\InfaSpark0",
\n2\n))), sqlContext.sql(InfaParams.resolve(InfaParams[String]("\InfaSpark0", \n3\n))));
\n Tuple2(sqlContext.sql(InfaParams.resolve(InfaParams[String]("\InfaSpark0",
\n4\n))), sqlContext.sql(InfaParams.resolve(InfaParams[String]("\InfaSpark0", \n5\n))));
\n val v0 = infabroadcast(asBlock(sqlContext.sql(InfaParams.resolve(InfaParams[String]
("\InfaSpark0", \n6\n")))).tag("\SRC_Read_students_5\n").itoDF("\m\n")).itoDF;\n val v1
= updatePartitions(asBlock(sqlContext.sql(InfaParams.resolve(InfaParams[String]
("\InfaSpark0", \n7\n")))).tag("\SRC_Read_students_HDFS_src\n").itoDF("\d\n"), v0);\n
val v2 = v1.join(v0, v0(0).==(v1(0)), \ninner\n).itoDF("\m\n");\n val v3 =
updatePartitions(asBlock(sqlContext.sql(InfaParams.resolve(InfaParams[String]
("\InfaSpark0", \n8\n")))).tag("\SRC_Read_student\n").itoDF("\d\n"), v2);\n val v4 =
v3.join(v2, v2(1).==(v3(0)), \ninner\n).itoDF;\n val v5 =
DataTypes.createDecimalType(28, 0);\n val v6 = DataTypes.createDecimalType(18, 0);
\n asBlock(sqlContext.sql(InfaParams.resolve(InfaParams[String]("\InfaSpark0",
\n9\n")))).v4.iselect(icast(icast(v4(2), v5), v6), v4(3), v4(4), icast(icast(v4(5), v5),
v6)).itoDF("\TGT_\n").tag("\TGT_Write_HDFSAppendTarget\n").itoDF("\c
\n").createOrReplaceTempView("\tbl0\n");\n } finally {\n
sqlContext.sql(InfaParams.resolve(InfaParams[String]("\InfaSpark0", \n10\n)));\n
sqlContext.sql(InfaParams.resolve(InfaParams[String]("\InfaSpark0", \n11\n)));\n
sqlContext.sql(InfaParams.resolve(InfaParams[String]("\InfaSpark0", \n12\n)));\n }
\n sc.stop\n}\n\n[0] -> [DROP TABLE IF EXISTS
`default`.`w7939778750618549156_infa_read_students_hdfs_src_hdfstgtappend_multipartition_
sparkmode`\n[1] -> [CREATE TABLE
`default`.`w7939778750618549156_infa_read_students_hdfs_src_hdfstgtappend_multipartition_
sparkmode` (`col0` INT, `col1` STRING, `col2` STRING, `col3` INT, `col4` STRING) ROW
FORMAT SERDE 'com.informatica.platform.dtm.executor.hive.boot.storagehandler.INFASerde'
STORED AS INPUTFORMAT
'com.informatica.platform.dtm.executor.hive.boot.storagehandler.INFAInputFormat'
OUTPUTFORMAT
'com.informatica.platform.dtm.executor.hive.boot.storagehandler.INFAOutputFormat'
LOCATION 'hdfs://nameservice1/tmp/SPARK_impUser1/sess7939778750618549156//
W7939778750618549156_infa_Read_students_HDFS_src_HDFSTgtAppend_MultiPartition_SparkMode'
TBLPROPERTIES ('infa.columns.types'='int,string,string,int,string',
'pwx.mapping.file.path'='./
Read_students_HDFS_src_MAPPING_37960411407997671_37960411786739094.bin',
'auto.purge'='true', 'infa.columns'='col0,col1,col2,col3,col4')]\n[2] -> [DROP TABLE IF
EXISTS
`default`.`w7939778750618549156_infa_read_students_5_hdfstgtappend_multipartition_sparkmo
de`\n[3] -> [CREATE TABLE
`default`.`w7939778750618549156_infa_read_students_5_hdfstgtappend_multipartition_sparkmo
de` (`col0` INT, `col1` STRING, `col2` STRING, `col3` INT) ROW FORMAT SERDE
'com.informatica.platform.dtm.executor.hive.boot.storagehandler.INFASerde' STORED AS
INPUTFORMAT
'com.informatica.platform.dtm.executor.hive.boot.storagehandler.INFAInputFormat'
OUTPUTFORMAT
'com.informatica.platform.dtm.executor.hive.boot.storagehandler.INFAOutputFormat'
LOCATION 'hdfs://nameservice1/tmp/SPARK_impUser1/sess7939778750618549156//
W7939778750618549156_INFACOPY_Read_students_5_HDFSTgtAppend_MultiPartition_SparkMode'
TBLPROPERTIES ('infa.columns.types'='int,string,string,int', 'pwx.mapping.file.path'='./
Read_students_5_MAPPING_37960411392603831_37960411887963169.bin', 'auto.purge'='true',
'infa.columns'='col0,col1,col2,col3')]\n[4] -> [DROP TABLE IF EXISTS
`default`.`w7939778750618549156_infa_write_hdfsappendtarget_hdfstgtappend_multipartition_
sparkmode`\n[5] -> [CREATE TABLE
`default`.`w7939778750618549156_infa_write_hdfsappendtarget_hdfstgtappend_multipartition_
sparkmode` (`col0` DECIMAL(18, 0), `col1` STRING, `col2` STRING, `col3` DECIMAL(18, 0))
ROW FORMAT SERDE
'com.informatica.platform.dtm.executor.hive.boot.storagehandler.INFASerde' STORED AS
INPUTFORMAT
'com.informatica.platform.dtm.executor.hive.boot.storagehandler.INFAInputFormat'
OUTPUTFORMAT
'com.informatica.platform.dtm.executor.hive.boot.storagehandler.INFAOutputFormat'
LOCATION 'hdfs://nameservice1/tmp/SPARK_impUser1/sess7939778750618549156//
W7939778750618549156_infa_Write_HDFSAppendTarget_HDFSTgtAppend_MultiPartition_SparkMode'
TBLPROPERTIES ('infa.columns.types'='decimal(18,0),string,string,decimal(18,0)',

```

```

'pwx.mapping.file.path'='./
Write_HDFSAppendTarget_MAPPING_37960411526174778_37960411903682194.bin',
'pwx.skip.serialization'='true', 'auto.purge'='true',
'infra.columns'='col0,col1,col2,col3']\n[6] -> [SELECT
`w7939778750618549156_infa_read_students_5_hdfstgtappend_multipartition_sparkmode`.`col0`
as a0,
`w7939778750618549156_infa_read_students_5_hdfstgtappend_multipartition_sparkmode`.`col1`
as a1,
`w7939778750618549156_infa_read_students_5_hdfstgtappend_multipartition_sparkmode`.`col2`
as a2,
`w7939778750618549156_infa_read_students_5_hdfstgtappend_multipartition_sparkmode`.`col3`
as a3 FROM
`default`.`w7939778750618549156_infa_read_students_5_hdfstgtappend_multipartition_sparkmo
de`]\n[7] -> [SELECT
`w7939778750618549156_infa_read_students_hdfs_src_hdfstgtappend_multipartition_sparkmode`
.`col0` as a0 FROM
`default`.`w7939778750618549156_infa_read_students_hdfs_src_hdfstgtappend_multipartition_
sparkmode`]\n[8] -> [SELECT alias.id as a0 FROM DEFAULT.student alias]\n[9] -> [INSERT
OVERWRITE TABLE
`default`.`w7939778750618549156_infa_write_hdfsappendtarget_hdfstgtappend_multipartition_
sparkmode` SELECT tbl0.c0 as a0, tbl0.c1 as a1, tbl0.c2 as a2, tbl0.c3 as a3 FROM
tbl0]\n[10] -> [DROP TABLE IF EXISTS
`default`.`w7939778750618549156_infa_write_hdfsappendtarget_hdfstgtappend_multipartition_
sparkmode`]\n[11] -> [DROP TABLE IF EXISTS
`default`.`w7939778750618549156_infa_read_students_5_hdfstgtappend_multipartition_sparkmo
de`]\n[12] -> [DROP TABLE IF EXISTS
`default`.`w7939778750618549156_infa_read_students_hdfs_src_hdfstgtappend_multipartition_
sparkmode`]"]
}
]
}

```

MappingExecutionPlans

With the MappingExecutionPlans REST API, you can view the execution plan for Hadoop jobs.

GET Request

To request information from the MappingExecutionPlans, use the following URL:

```
<RESTOperationsHubService_Host>:<RESTOperationsHubService_Port>/RestOperationsHub/
services/v1/MappingService/MappingExecutionPlans('jobId')
```

The following table describes the required attributes in the MappingExecutionPlans Get URL:

Field	Type	Description
userName	String	Required. User name to connect to the domain. You can pass the input value as a header.
encryptedpassword	String	Required. Password for the user. Encrypt the password with the pmpasswd command line program. You can pass the input value as a header.

Field	Type	Description
securityDomain	String	Optional. The security domain to which the domain user belongs. You can pass the input value as a header.
jobId	String	Required. The argument of the entity that contains the ID for the mappings. You can pass the input value as a query.

Get Response

Return information for the MappingExecutionPlans for the specified Job ID.

The following table describes the MappingExecutionPlans attributes present in the body of the response for the Spark or Blaze environment:

Field	Type	Description
name	Integer	Name of the script.
content	String	For Spark, the Data Integration Service translates the mapping to a Scala program and an optional set of commands. The execution plan shows the commands and the Scala program code. For Blaze, the content comprises of session task, instances, and type.
depends	String	Tasks that the script depends on. Tasks include other scripts and Data Integration Service tasks, like the Start task.

The following table describes the MappingExecutionPlans attributes present in the content section of the body of the response for Blaze environment:

Field	Description
Session Task	Response can include three types of session tasks that are pre session task, main session task, and post session task. The main session task can contain Submappings.
Type	Type of the session task containing a set of segments and DAG vertices.
Instances	Transformation or object name.

Sample Retrieve Mapping Execution Plans

The sample use case is to use the script to retrieve the details of the mapping execution plans.

You can use the REST API to retrieve information about the mapping execution plans with the following request URL for a mapping with Job ID as `_TNo09ELEeimY76kFyfuw`:

```
<RESTOperationsHubService_Host>:<RESTOperationsHubService_Port>/restopshub/services/v1/
MappingService/MappingExecutionPlans('_TNo09ELEeimY76kFyfuw')
```

Mapping Execution Plans Output

```
{
  "@odata.context": "$metadata#MappingExecutionPlans/$entity",
  "jobId": "_TNo09ELEeimY76kFyfuw",
  "scriptList": [
    {
```

```

        "name": "InfaSpark0",
        "content": "package com.informatica.exec\n\nimport
com.informatica.bootstrap.functions._\nimport com.informatica.bootstrap.InfaParams._
\nimport com.informatica.bootstrap.InfaStreaming.writeToKafka\nimport
com.informatica.products.infatransform.spark.boot._\nimport com.informatica.bootstrap._
\nimport com.informatica.hive._\nimport com.informatica.bootstrap.{JsonProtocol => JP}\n
\nimport org.apache.spark._\nimport org.apache.spark.rdd._\nimport
org.apache.spark.storage.StorageLevel._\nimport org.apache.spark.sql._\nimport
org.apache.spark.sql.types._\nimport org.apache.spark.sql.functions._\nimport
org.apache.spark.sql.functions.{ broadcast => infabroadcast }\nimport
org.apache.spark.sql.infa.expressions._\nimport java.io._\nimport java.sql.Timestamp
\nimport scala.reflect.ClassTag\nimport org.apache.spark.sql.catalyst.expressions.Caster
\nimport org.apache.spark.sql.catalyst.expressions.JavaCaster\nimport
com.informatica.bootstrap.JavaTx._\nimport org.apache.spark.Accumulator\nimport
org.apache.spark.util.LongAccumulator\nimport org.apache.spark.scheduler.SparkListener
\nimport org.apache.spark.SparkEnv\nimport org.apache.spark.sql.Row\n\nobject InfaSpark0
{\n  def main(s:Array[String]) {\n    val sc = SparkContextLoader.getSparkContext\n
val sqlContext = SparkContextLoader.getSQLContext\n    val ls = new
LiveStream(sc.getConf)\n    ls.relay(JP.sparkConfToJson(sc.getConf))\n    val lis = new
InfaListener(ls,\"TAG\")\n    sc.addSparkListener(lis)
\nsqlContext.sparkSession.experimental.extraPreprocessing = new InfaTaggingRules().rules
\n    val accs = List()\n    ls.relay(JP.sparkAppDetailsToJson(sc.getConf, accs))\n
lis.accumulators = accs\n    import sqlContext.implicits._\n    import
org.apache.spark.sql.functions.{stddev_samp, var_samp}\n    val icast = caster(\"MM/DD/
YYYY HH24:MI:SS\")\n    val acast = adapterCaster()\n    val jcast = JavaCaster()\n\n
try {\n    Tuple2(sqlContext.sql(InfaParams.resolve(InfaParams[String](\"InfaSpark0\",
\n0\"))), sqlContext.sql(InfaParams.resolve(InfaParams[String](\"InfaSpark0\", \"1\"))));\n
\n    Tuple2(sqlContext.sql(InfaParams.resolve(InfaParams[String](\"InfaSpark0\",
\n2\"))), sqlContext.sql(InfaParams.resolve(InfaParams[String](\"InfaSpark0\", \"3\"))));\n
\n    Tuple2(sqlContext.sql(InfaParams.resolve(InfaParams[String](\"InfaSpark0\",
\n4\"))), sqlContext.sql(InfaParams.resolve(InfaParams[String](\"InfaSpark0\", \"5\"))));\n
\n    val v0 = infabroadcast(asBlock(sqlContext.sql(InfaParams.resolve(InfaParams[String]
(\"InfaSpark0\", \"6\"))))).tag(\"SRC_Read_students_5\").itoDF(\"m\")).itoDF;\n    val v1
= updatePartitions(asBlock(sqlContext.sql(InfaParams.resolve(InfaParams[String]
(\"InfaSpark0\", \"7\"))))).tag(\"SRC_Read_students_HDFS_src\").itoDF(\"d\"), v0);\n
val v2 = v1.join(v0, v0(0).===(v1(0)), \"inner\").itoDF(\"m\");\n    val v3 =
updatePartitions(asBlock(sqlContext.sql(InfaParams.resolve(InfaParams[String]
(\"InfaSpark0\", \"8\"))))).tag(\"SRC_Read_student\").itoDF(\"d\"), v2);\n    val v4 =
v3.join(v2, v2(1).===(v3(0)), \"inner\").itoDF;\n    val v5 =
DataTypes.createDecimalType(28, 0);\n    val v6 = DataTypes.createDecimalType(18, 0);
\n    asBlock(sqlContext.sql(InfaParams.resolve(InfaParams[String](\"InfaSpark0\",
\n9\"))), v4.iselect(icast(icast(v4(2)), v5), v6), v4(3), v4(4), icast(icast(v4(5)), v5),
v6).itoDF(\"TGT_\").tag(\"TGT_Write_HDFSAppendTarget\").itoDF(\"c
\n\").createOrReplaceTempView(\"tbl0\");\n    } finally {\n
sqlContext.sql(InfaParams.resolve(InfaParams[String](\"InfaSpark0\", \"10\")));;\n
sqlContext.sql(InfaParams.resolve(InfaParams[String](\"InfaSpark0\", \"11\")));;\n
sqlContext.sql(InfaParams.resolve(InfaParams[String](\"InfaSpark0\", \"12\")));;\n    }
\n    sc.stop\n\n}\n\n[0] -> [DROP TABLE IF EXISTS
`default`.`w7939778750618549156_infa_read_students_hdfs_src_hdfstgtappend_multipartition_
sparkmode`]\n[1] -> [CREATE TABLE
`default`.`w7939778750618549156_infa_read_students_hdfs_src_hdfstgtappend_multipartition_
sparkmode` (`col0` INT, `col1` STRING, `col2` STRING, `col3` INT, `col4` STRING) ROW
FORMAT SERDE 'com.informatica.platform.dtm.executor.hive.boot.storagehandler.INFASerde'
STORED AS INPUTFORMAT
'com.informatica.platform.dtm.executor.hive.boot.storagehandler.INFAInputFormat'
OUTPUTFORMAT
'com.informatica.platform.dtm.executor.hive.boot.storagehandler.INFAOutputFormat'
LOCATION 'hdfs://nameservice1/tmp/SPARK_impUser1/sess7939778750618549156//
W7939778750618549156_infa_read_students_HDFS_src_HDFStgtAppend_MultiPartition_SparkMode'
TBLPROPERTIES ('infa.columns.types'='int,string,string,int,string',
'pwx.mapping.file.path'='./
Read_students_HDFS_src_MAPPING_37960411407997671_37960411786739094.bin',
'auto.purge'='true', 'infa.columns'='col0,col1,col2,col3,col4')]\n[2] -> [DROP TABLE IF
EXISTS
`default`.`w7939778750618549156_infa_read_students_5_hdfstgtappend_multipartition_
sparkmode`]\n[3] -> [CREATE TABLE
`default`.`w7939778750618549156_infa_read_students_5_hdfstgtappend_multipartition_
sparkmode` (`col0` INT, `col1` STRING, `col2` STRING, `col3` INT) ROW FORMAT SERDE
'com.informatica.platform.dtm.executor.hive.boot.storagehandler.INFASerde' STORED AS
INPUTFORMAT

```

```

'com.informatica.platform.dtm.executor.hive.boot.storagehandler.INFAInputFormat'
OUTPUTFORMAT
'com.informatica.platform.dtm.executor.hive.boot.storagehandler.INFAOutputFormat'
LOCATION 'hdfs://nameservice1/tmp/SPARK_impUser1/sess7939778750618549156//
W7939778750618549156_INFACOPY_Read_students_5_HDFSTgtAppend_MultiPartition_SparkMode'
TBLPROPERTIES ('infa.columns.types'='int,string,string,int', 'pwx.mapping.file.path'='./
Read_students_5_MAPPING_37960411392603831_37960411887963169.bin', 'auto.purge'='true',
'infa.columns'='col0,col1,col2,col3']\n[4] -> [DROP TABLE IF EXISTS
`default`.w7939778750618549156_infa_write_hdfsappendtarget_hdfstgtappend_multipartition_
sparkmode`\n[5] -> [CREATE TABLE
`default`.w7939778750618549156_infa_write_hdfsappendtarget_hdfstgtappend_multipartition_
sparkmode` (`col0` DECIMAL(18, 0), `col1` STRING, `col2` STRING, `col3` DECIMAL(18, 0))
ROW FORMAT SERDE
'com.informatica.platform.dtm.executor.hive.boot.storagehandler.INFASerde' STORED AS
INPUTFORMAT
'com.informatica.platform.dtm.executor.hive.boot.storagehandler.INFAInputFormat'
OUTPUTFORMAT
'com.informatica.platform.dtm.executor.hive.boot.storagehandler.INFAOutputFormat'
LOCATION 'hdfs://nameservice1/tmp/SPARK_impUser1/sess7939778750618549156//
W7939778750618549156_infa_Write_HDFSAppendTarget_HDFSTgtAppend_MultiPartition_SparkMode'
TBLPROPERTIES ('infa.columns.types'='decimal(18,0),string,string,decimal(18,0)',
'pwx.mapping.file.path'='./
Write_HDFSAppendTarget_MAPPING_37960411526174778_37960411903682194.bin',
'pwx.skip.serialization'='true', 'auto.purge'='true',
'infa.columns'='col0,col1,col2,col3']\n[6] -> [SELECT
`w7939778750618549156_infa_read_students_5_hdfstgtappend_multipartition_sparkmode`.`col0`
as a0,
`w7939778750618549156_infa_read_students_5_hdfstgtappend_multipartition_sparkmode`.`col1`
as a1,
`w7939778750618549156_infa_read_students_5_hdfstgtappend_multipartition_sparkmode`.`col2`
as a2,
`w7939778750618549156_infa_read_students_5_hdfstgtappend_multipartition_sparkmode`.`col3`
as a3 FROM
`default`.w7939778750618549156_infa_read_students_5_hdfstgtappend_multipartition_sparkmo
de`\n[7] -> [SELECT
`w7939778750618549156_infa_read_students_hdfs_src_hdfstgtappend_multipartition_sparkmode`
.`col0` as a0 FROM
`default`.w7939778750618549156_infa_read_students_hdfs_src_hdfstgtappend_multipartition_
sparkmode`\n[8] -> [SELECT alias.id as a0 FROM DEFAULT.student alias]\n[9] -> [INSERT
OVERWRITE TABLE
`default`.w7939778750618549156_infa_write_hdfsappendtarget_hdfstgtappend_multipartition_
sparkmode` SELECT tbl0.c0 as a0, tbl0.c1 as a1, tbl0.c2 as a2, tbl0.c3 as a3 FROM
tbl0]\n[10] -> [DROP TABLE IF EXISTS
`default`.w7939778750618549156_infa_write_hdfsappendtarget_hdfstgtappend_multipartition_
sparkmode`\n[11] -> [DROP TABLE IF EXISTS
`default`.w7939778750618549156_infa_read_students_5_hdfstgtappend_multipartition_sparkmo
de`\n[12] -> [DROP TABLE IF EXISTS
`default`.w7939778750618549156_infa_read_students_hdfs_src_hdfstgtappend_multipartition_
sparkmode`]`,
  "depends": "Pre_Spark_Task_Command_1"
},
{
  "name": "Pre_Spark_Task_Command_1",
  "content": "-- ----\n-- Command [Pre_Spark_Task_Command_1_1] \n/data/
Informatica/10.2.2_252/services/shared/hadoop/HDP_2.6/scripts/FileCopyUtil --
hadoop.home /data/Informatica/10.2.2_252/services/shared/hadoop/HDP_2.6 --hdfsUser
impUser1 --copyFromLocal --spn adpqa@INFAKRB.INFADEV.COM --keytab /bdmqa/BDM_Automation/
Source/adpqa_AD.keytab --ccoConfPath /data/Informatica/10.2.2_252/tomcat/bin/disTemp/
inkrh71hdp07_252/DIS_HDP_2.6/cco_hdp_26/SPARK/infacco-site.xml file:///bdmqa/
BDM_Automation/Source/students_10.txt hdfs://nameservice1/tmp/SPARK_impUser1/
sess7939778750618549156//
W7939778750618549156_INFACOPY_Read_students_5_HDFSTgtAppend_MultiPartition_SparkMode/
students_10.txt\n-- ----\n-- Command [Pre_Spark_Task_Command_1_2] \n/data/Informatica/
10.2.2_252/services/shared/hadoop/HDP_2.6/scripts/HadoopFsMkdir --hadoop.home /data/
Informatica/10.2.2_252/services/shared/hadoop/HDP_2.6 --hdfsUser impUser1 --spn
adpqa@INFAKRB.INFADEV.COM --keytab /bdmqa/BDM_Automation/Source/adpqa_AD.keytab --
ccoConfPath /data/Informatica/10.2.2_252/tomcat/bin/disTemp/inkrh71hdp07_252/
DIS_HDP_2.6/cco_hdp_26/SPARK/infacco-site.xml hdfs://nameservice1/tmp/SPARK_impUser1/
sess7939778750618549156//
W7939778750618549156_Write_HDFSAppendTarget_HDFSTgtAppend_MultiPartition_SparkMode/
reject-files\n-- ----\n-- Command [Pre_Spark_Task_Command_1_3] \n/data/Informatica/

```



```

10.2.2_252/services/shared/hadoop/HDP_2.6/scripts/HadoopFsRmRf --hadoop.home /data/
Informatica/10.2.2_252/services/shared/hadoop/HDP_2.6 --hdfsUser impUser1 --spn
adpqa@INFAKRB.INFADEV.COM --keytab /bmq/BDM_Automation/Source/adpqa_AD.keytab --
ccoConfPath /data/Informatica/10.2.2_252/tomcat/bin/disTemp/inkrh71hdp07_252/
DIS_HDP_2.6/ccco_hdp_26/SPARK/infacco-site.xml hdfs://nameservice1/BDM_Automation/Target/
HDFSTargetAppend/SparkMode/MultiPartition//StudentHDFSTargetAppend.out-[mr]?[0-9]*\n--
----\n-- Command [Pre Spark Task Command 1 4] \ntouch /bmq/BDM_Automation/Target/
StudentHDFSTargetAppend61c555e1-c967-4888-bc7d-ab9ada8ee2a7_empty.txt\n-- ----\n--
Command [Pre Spark Task Command 1 5] \n/data/Informatica/10.2.2_252/services/shared/
hadoop/HDP_2.6/scripts/FileCopyUtil --hadoop.home /data/Informatica/10.2.2_252/services/
shared/hadoop/HDP_2.6 --hdfsUser impUser1 --copyFromLocal --spn
adpqa@INFAKRB.INFADEV.COM --keytab /bmq/BDM_Automation/Source/adpqa_AD.keytab --
ccoConfPath /data/Informatica/10.2.2_252/tomcat/bin/disTemp/inkrh71hdp07_252/
DIS_HDP_2.6/ccco_hdp_26/SPARK/infacco-site.xml file:///bmq/BDM_Automation/Target/
StudentHDFSTargetAppend61c555e1-c967-4888-bc7d-ab9ada8ee2a7_empty.txt hdfs://
nameservice1/BDM_Automation/Target/HDFSTargetAppend/SparkMode/MultiPartition//
StudentHDFSTargetAppend.out\n-- ----\n-- Command [Pre Spark Task Command 1 6] \n/data/
Informatica/10.2.2_252/services/shared/hadoop/HDP_2.6/scripts/HadoopFsRmRf --
hadoop.home /data/Informatica/10.2.2_252/services/shared/hadoop/HDP_2.6 --hdfsUser
impUser1 --spn adpqa@INFAKRB.INFADEV.COM --keytab /bmq/BDM_Automation/Source/
adpqa_AD.keytab --ccoConfPath /data/Informatica/10.2.2_252/tomcat/bin/disTemp/
inkrh71hdp07_252/DIS_HDP_2.6/ccco_hdp_26/SPARK/infacco-site.xml hdfs://nameservice1/
BDM_Automation/Target/HDFSTargetAppend/SparkMode/MultiPartition//
*_tmp_infa_7939778750618549156_StudentHDFSTargetAppend.out\n-- ----\n-- Command
[Pre Spark Task Command 1 7] \n/data/Informatica/10.2.2_252/services/shared/hadoop/
HDP_2.6/scripts/HadoopFsMkdir --hadoop.home /data/Informatica/10.2.2_252/services/shared/
hadoop/HDP_2.6 --hdfsUser impUser1 --spn adpqa@INFAKRB.INFADEV.COM --keytab /bmq/
BDM_Automation/Source/adpqa_AD.keytab --ccoConfPath /data/Informatica/10.2.2_252/
tomcat/bin/disTemp/inkrh71hdp07_252/DIS_HDP_2.6/ccco_hdp_26/SPARK/infacco-site.xml hdfs://
nameservice1/BDM_Automation/Target/HDFSTargetAppend/SparkMode/MultiPartition//
*_tmp_infa_7939778750618549156_StudentHDFSTargetAppend.out/\n",
  "depends": ""
},
{
  "name": "Post Spark Task Command 1",
  "content": "-- ----\n-- Command [Post Spark Task Command 1 1] \n/data/
Informatica/10.2.2_252/services/shared/hadoop/HDP_2.6/scripts/FileCopyUtil --
hadoop.home /data/Informatica/10.2.2_252/services/shared/hadoop/HDP_2.6 --mergefiles --
deleteEmptyFiles --hdfsUser impUser1 --copyToLocal --spn adpqa@INFAKRB.INFADEV.COM --
keytab /bmq/BDM_Automation/Source/adpqa_AD.keytab --ccoConfPath /data/Informatica/
10.2.2_252/tomcat/bin/disTemp/inkrh71hdp07_252/DIS_HDP_2.6/ccco_hdp_26/SPARK/infacco-
site.xml hdfs://nameservice1//tmp/SPARK_impUser1/sess7939778750618549156//
W7939778750618549156_Write_HDFSAppendTarget_HDFSTgtAppend_MultiPartition_SparkMode/
reject-files file:///data/Informatica/10.2.2_252/tomcat/bin/reject/.bad\n",
  "depends": "InfaSpark0"
}
]
}

```

INDEX

A

- active properties view [44](#)
- ADLS certificate contents [66](#)
- Amazon AWS [61](#), [83](#)
- Amazon Redshift connection properties [88](#)
- Amazon S3 connection properties [89](#)
- Apache Knox [22](#)
- Apache Knox Gateway authentication [24](#)
- architecture
 - Big Data Management with Databricks [15](#)
 - Hadoop environment [11](#)
- authentication
 - Apache Knox Gateway [24](#)
 - infrastructure security [21](#)
 - Kerberos [23](#)
- authentication systems
 - Knox [22](#)
 - LDAP [22](#)
 - SASL [22](#)
- authorization
 - HDFS permissions [36](#)
 - HDFS Transparent Encryption [36](#)
 - infrastructure security [35](#)
 - Ranger [36](#), [37](#)
 - Sentry [36](#)
 - SQL authorization [37](#)
- autotune
 - connections [78](#)
 - domain [78](#)
 - services [78](#)
- Azure
 - configuration [63](#), [84](#)

B

- big data
 - application services [16](#)
- Blaze
 - timeout [14](#)
- Blaze engine
 - Blaze engine architecture [13](#)
 - connection properties [97](#)

C

- Cassandra connections
 - properties [91](#)
- cloud provisioning configuration
 - Databricks properties [66](#), [87](#)
 - Amazon AWS properties [61](#), [83](#)

- cloud provisioning configuration (*continued*)
 - Microsoft Azure properties [63](#), [84](#)
- cluster configuration
 - active properties [44](#)
 - create user-defined properties [54](#)
 - import prerequisites [46](#)
 - overriding properties [53](#)
- connections [43](#)
 - create [46](#)
 - deleting [57](#)
 - editing [52](#)
 - import [46](#)
 - import from a cluster [47](#)
 - import from a file [48](#)
 - permission types [58](#)
 - properties
 - creating [54](#)
 - deleting [55](#)
 - overridden [45](#)
 - refresh [56](#)
 - views [44](#)
- cluster configuration properties
 - filtering [52](#)
- cluster information
 - prerequisite [46](#)
- cluster workflow
 - cloud provisioning connection [82](#)
- configuration sets [42](#), [44](#), [52](#), [59](#), [66](#)
- connecting to a cluster [47](#)
- Connection
 - details [114](#)
 - properties [114](#)
- connection properties
 - Databricks [92](#)
- connections
 - properties [82](#), [97](#)
 - HBase [82](#)
 - HDFS [82](#)
 - Hive [82](#)
 - JDBC [82](#)
 - to an SSL-enabled cluster [39](#)
- Cosmos DB connection
 - creating [115](#)
- create cluster configuration [46](#)
- creating
 - Cosmos DB connection [115](#)
- cross-realm trust
 - Kerberos authentication [28](#)

D

- Data Lake Service Principal Certificate Contents [66](#)
- data management
 - HDFS Transparent Encryption [36](#)
 - Sentry [36](#)

- Databricks
 - cloud provisioning configuration [66, 87](#)
 - import file [50](#)
 - import from file [51](#)
 - import from cluster [49](#)
- Databricks connection properties [92](#)
- Databricks integration
 - overview [15](#)
- delete cluster configuration [57](#)
- deleting properties [55](#)
- deployment type
 - advanced [75](#)
 - basic [75](#)
 - Hadoop [75](#)
 - sandbox [75](#)
 - standard [75](#)
- domain objects
 - cloud provisioning configuration [59, 66](#)
 - cluster configuration [42](#)

E

- edit cluster configuration [52](#)
- ephemeral clusters
 - cloud provisioning connection [82](#)
- example
 - cluster configuration refresh [56](#)

F

- filter
 - cluster configuration properties [52](#)

G

- Google Analytics connections
 - properties [94](#)
- Google BigQuery connection
 - properties [94](#)
- Google Cloud Spanner connection
 - properties [95](#)
- Google Cloud Storage connections
 - properties [96](#)
- Grid Manager [127](#)

H

- Hadoop [82](#)
- Hadoop connections
 - creating [119](#)
- hadoop utilities
 - Sqoop [11](#)
- HBase connections
 - MapR-DB properties [104](#)
 - properties [104](#)
- HDFS connections
 - creating [119](#)
 - properties [102](#)
- HDFS permissions
 - authorization [36](#)
- HDFS Transparent Encryption
 - authorization [36](#)
 - data management [36](#)

- Hive
 - authorization [37](#)
- Hive connections
 - creating [119](#)
 - properties [105](#)
- Hive pushdown
 - connection properties [97](#)

I

- import file
 - Databricks [50](#)
- import from cluster
 - Databricks [49](#)
- import from file
 - Databricks [51](#)
- infacmd autotune
 - Autotune [78](#)
- infagrid.blaze.service.idle.timeout [14](#)
- infagrid.orchestrator.svc.sunset.time [14](#)
- infrastructure security
 - authentication [21](#)
 - authorization [35](#)

J

- JDBC connections
 - properties [108](#)
- job recovery [68, 72](#)

K

- Kerberos authentication
 - cross-realm trust [28](#)
 - mappings in a native environment [34](#)
 - overview [25](#)
 - user impersonation [32](#)
 - user impersonation in the native environment [33](#)
- Knox [22](#)
- krb5.conf [26](#)

L

- LDAP [22](#)

M

- MappingAdvancedStats [140](#)
- MappingExecutionPlans [149](#)
- MappingExecutionSteps [145](#)
- MappingStats [138](#)
- Metadata Document [136](#)
- Microsoft Azure [63, 84](#)
- Microsoft Azure Data Lake Store connection
 - properties [116](#)
- Microsoft Azure SQL Data Warehouse connection
 - properties [117](#)
- Monitoring [135, 136](#)

O

- overridden properties view [44](#)

P

- permissions
 - definition [57](#)
- privileges and roles
 - administrator role [57](#)
 - domain administration privilege [57](#)
 - manage connections privilege [57](#)

R

- Ranger
 - authorization [36](#), [37](#)
- recovery [68](#), [72](#)
- refresh cluster configuration [56](#)
- REST API [135](#), [138](#), [140](#), [145](#), [149](#)

S

- SASL [22](#)
- security certificates [38](#), [39](#)
- Sentry
 - authorization [36](#)
 - data management [36](#)
- Snowflake connection
 - properties [118](#)
- Spark deploy mode
 - Hadoop connection properties [97](#)
- Spark engine
 - connection properties [97](#)
- Spark Event Log directory
 - Hadoop connection properties [97](#)
- Spark execution parameters
 - Hadoop connection properties [97](#)
- Spark HDFS staging directory
 - Hadoop connection properties [97](#)
- SQL authorization [36](#), [37](#)
- Sqoop connection arguments
 - Dsqoop.connection.factories [111](#)
 - connect [111](#)
 - direct [111](#)
 - driver [111](#)
- SSL [38](#), [39](#)

SSL security protocol [36](#)

T

- TDCH connection factory
 - Dsqoop.connection.factories [111](#)
- third-party tools
 - hadoop cluster [11](#)
- timeout
 - Blaze [14](#)
- TLS security protocol [36](#)
- tune
 - Blaze [77](#)
 - Content Management Service [75](#)
 - Data Integration Service [75](#)
 - engines [76](#)
 - Hadoop [75-77](#)
 - Model Repository Service [75](#)
 - Search Service [75](#)
 - services [75](#)
 - Spark [77](#)
- tuning
 - advanced [75](#)
 - basic [75](#)
 - engines [74](#)
 - Hadoop [74](#)
 - sandbox [75](#)
 - services [74](#)
 - standard [75](#)

U

- user impersonation
 - Hadoop environment [32](#)
- user-defined properties
 - creating [54](#)
- utilities
 - hadoop cluster [11](#)

V

views [44](#)